

Cosmic dynamics in the era of Extremely Large Telescopes

J. Liske,^{1*} A. Grazian,² E. Vanzella,³ M. Dessauges,⁴ M. Viel,^{3,5} L. Pasquini,¹
M. Haehnelt,⁵ S. Cristiani,³ F. Pepe,⁴ G. Avila,¹ P. Bonifacio,^{6,3} F. Bouchy,^{7,8}
H. Dekker,¹ B. Delabre,¹ S. D’Odorico,¹ V. D’Odorico,³ S. Levshakov,⁹ C. Lovis,⁴
M. Mayor,⁴ P. Molaro,³ L. Moscardini,^{10,11} M.T. Murphy,^{5,12} D. Queloz,⁴ P. Shaver,¹
S. Udry,⁴ T. Wiklind^{13,14} and S. Zucker¹⁵

¹European Southern Observatory, Karl-Schwarzschild-Str. 2, 85748 Garching, Germany

²INAF – Osservatorio Astronomico di Roma, via di Frascati 33, 00040 Monteporzio Catone (Roma), Italy

³INAF – Osservatorio Astronomico di Trieste, Via Tiepolo 11, 34143 Trieste, Italy

⁴Observatoire de Genève, 51 Ch. des Maillettes, 1290 Sauverny, Switzerland

⁵Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA

⁶CIFIST Marie Curie Excellence Team, GEPI, Observatoire de Paris, CNRS, Université Paris Diderot, Place Jules Janssen 92190 Meudon, France

⁷Laboratoire d’Astrophysique de Marseille, Traverse du Siphon, 13013 Marseille, France

⁸Observatoire de Haute-Provence, 04870 St Michel l’Observatoire, France

⁹Department of Theoretical Astrophysics, Ioffe Physico-Technical Institute, 194021 St. Petersburg, Russia

¹⁰Dipartimento di Astronomia, Università di Bologna, via Ranzani 1, 40127 Bologna, Italy

¹¹INFN – National Institute for Nuclear Physics, Sezione di Bologna, viale Bertini Pichat 6/2, 40127 Bologna, Italy

¹²Centre for Astrophysics & Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia

¹³Space Telescope Science Institute, 3700 San Martin Drive, Baltimore MD 21218, USA

¹⁴Affiliated with the Space Sciences Department of the European Space Agency

¹⁵Department of Geophysics and Planetary Sciences, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 69978, Israel

Accepted Received

ABSTRACT

The redshifts of all cosmologically distant sources are expected to experience a small, systematic drift as a function of time due to the evolution of the Universe’s expansion rate. A measurement of this effect would represent a direct and entirely model-independent determination of the expansion history of the Universe over a redshift range that is inaccessible to other methods. Here we investigate the impact of the next generation of Extremely Large Telescopes on the feasibility of detecting and characterising the cosmological redshift drift. We consider the Lyman α forest in the redshift range $2 < z < 5$ and other absorption lines in the spectra of high redshift QSOs as the most suitable targets for a redshift drift experiment. Assuming photon-noise limited observations and using extensive Monte Carlo simulations we determine the accuracy to which the redshift drift can be measured from the Ly α forest as a function of signal-to-noise and redshift. Based on this relation and using the brightness and redshift distributions of known QSOs we find that a 42-m telescope is capable of unambiguously detecting the redshift drift over a period of ~ 20 yr using 4000 h of observing time. Such an experiment would provide independent evidence for the existence of dark energy without assuming spatial flatness, using any other cosmological constraints or making any other astrophysical assumption.

Key words: cosmology: miscellaneous – intergalactic medium – quasars: absorption lines.

1 INTRODUCTION

The universal expansion was the first observational evidence that general relativity might be applicable to the Universe as a whole. Since Hubble’s (1929) discovery much effort has been invested into completing the basic picture of relativistic cosmology. The central

question is: what is the stress-energy tensor of the Universe? Assuming homogeneity and isotropy reduces this question to: what is the mean density and equation of state of each mass-energy component of the Universe? Since these parameters determine both the evolution with time and the geometry of the metric that solves the Einstein equation, one can use a measurement of either to infer their values. Over the past decade the successes on this front have reached their (temporary) culmination: observations of the

* E-mail: jliske@eso.org

Cosmic Microwave Background (CMB; Spergel et al. 2003, 2007), type Ia supernovae (SNIa; Riess et al. 2004; Astier et al. 2006), the large-scale galaxy distribution (Peacock et al. 2001; Cole et al. 2005; Eisenstein et al. 2005) and others now provide answers of such convincing consistency and accuracy that the term ‘precision cosmology’ is now commonplace (e.g. Primack 2005).

By far the most unexpected result of this campaign was the discovery that the expansion of the Universe has recently begun accelerating (Riess et al. 1998; Perlmutter et al. 1999). The physical reason for this acceleration is entirely unclear at present. Within relativistic cosmology it can be accommodated by modifying the stress-energy tensor to include a new component with negative pressure. In its simplest incarnation this so-called dark energy is the cosmological constant Λ (e.g. Carroll, Press & Turner 1992), i.e. a smooth, non-varying component with equation of state parameter $w = -1$. It could also be time variable (e.g. Overduin & Cooperstock 1998), either because of a specific equation of state (xCDM, e.g. Turner & White 1997; phantom energy, Caldwell 2002) or because the component consists of a dynamical scalar field evolving in a potential (e.g. Ratra & Peebles 1988). More generally it may be an inhomogeneous, time varying component with $-1 \lesssim w \lesssim 0$ (where w may also vary with time), sometimes called quintessence (e.g. Caldwell, Dave & Steinhardt 1998), or even a component with an exotic equation of state (e.g. Chaplygin gas, Kamenshchik, Moschella & Pasquier 2001). Alternatively, instead of modifying the stress-energy tensor one can also modify gravity itself to explain the acceleration (e.g. Deffayet, Dvali & Gabadadze 2002; Freese & Lewis 2002; Carroll et al. 2005). Although all current observations are consistent with a cosmological constant (e.g. Seljak, Slosar & McDonald 2006) many more unfamiliar models are not ruled out (e.g. Capozziello et al. 2005).

Probably the best way to probe the nature of the acceleration is to determine the expansion history of the Universe (Linder 2003; Seo & Eisenstein 2003). Observables that depend on the expansion history include distances and the linear growth of density perturbations (Linder & Jenkins 2003; Linder 2005), and so SNIa surveys, weak lensing (Heavens 2003; Jain & Taylor 2003) and baryon acoustic oscillations in the galaxy power spectrum (BAO; Seo & Eisenstein 2003; Wang 2006) are generally considered to be excellent probes of the acceleration.

In practice, however, extracting information on the expansion history from weak lensing and BAO requires a prior on the spatial curvature, a detailed understanding of the linear growth of density perturbations and hence a specific cosmological model. Given the uncertain state of affairs regarding the source of the acceleration, and given that even simple parameterisations of dark energy properties can result in misleading conclusions (Maor et al. 2001; Bassett et al. 2004), these are conceptually undesirable features and several authors have pointed out the importance of taking a cosmographic, model-independent approach to determining the expansion history (e.g. Wang & Tegmark 2005; John 2005; Shapiro & Turner 2006). Using SNIa to measure luminosity distances as a function of redshift is conceptually the simplest experiment and hence appears to be the most useful in this respect. The caveats are that distance is ‘only’ related to the expansion history through an integral over redshift and that one still requires a prior on spatial curvature (Caldwell & Kamionkowski 2004).

Here we will revisit a method to directly measure the expansion history that was first explored by Sandage (1962, with an appendix by McVittie). He showed that the evolution of the Hubble expansion causes the redshifts of distant objects partaking in the

Hubble flow to change slowly with time. Just as the redshift, z , is in itself evidence of the expansion, so is the change in redshift, \dot{z} , evidence of its de- or acceleration between the epoch z and today. This implies that the expansion history can be determined, at least in principle, by means of a straightforward spectroscopic monitoring campaign.

The redshift drift is a direct, entirely model-independent measurement of the expansion history of the Universe which does not require any cosmological assumptions or priors whatsoever. However, the most unique feature of this experiment is that it directly probes the global *dynamics* of the metric. All other existing cosmological observations, including those of the CMB, SNIa, weak lensing and BAO, are essentially *geometric* in nature in the sense that they map out space, its curvature and its evolution. Many of these experiments also probe the dynamics of localised density perturbations but none actually measure the *global* dynamics. In this sense the redshift drift method is qualitatively different from all other cosmological observations, offering a truly independent and unique approach to the exploration of the expansion history of the Universe.

Following the original study by Sandage (1962), the redshift drift and its relevance to observational cosmology were also discussed by McVittie (1965), Weinberg (1972), Ebert & Trümper (1975), Rüdiger (1980), Peacock (1999), Nakamura & Chiba (1999), Zhu & Fujimoto (2004), Corasaniti, Huterer & Melchiorri (2007) and Lake (2007). Lake (1981) gave equations expressing the deceleration and matter density parameters, q_0 and Ω_M , in terms of z , \dot{z} and \ddot{z} . An excellent exposé of the equations relevant to redshift evolution was also presented by Gudmundsson & Björnsson (2002). In addition, these authors investigated the redshift drift in the presence of quintessence, while other non-standard dark energy models were considered by Balbi & Quercellini (2007) and Zhang et al. (2007). The case of Dicke-Brans-Jordan cosmologies was scrutinised by Rüdiger (1982) and Partovi & Mashhoon (1984) studied the effects of inhomogeneities. Seto & Cooray (2006) suggested that a measurement of the anisotropy of \dot{z} could be used to constrain the very low-frequency gravitational wave background. Several of these authors have pointed out the superior redshift accuracy achieved in the radio regime compared to the optical, and Davis & May (1978) entertained the possibility of using a 21 cm absorption line at $z = 0.69$ in the radio spectrum of 3C 286 for a \dot{z} measurement. The detrimental effects of peculiar velocities and accelerations, which may swamp the cosmic signal, were discussed by Phillipps (1982), Lake (1982) and Teuber (1986). Finally, Loeb (1998) first proposed the Lyman α (Ly α) forest as an appropriate target for a \dot{z} measurement (acknowledging D. Sasselov for the idea) and assessed the prospects for a successful detection in the context of currently existing observational technology. All except the last of these studies concluded that a \dot{z} measurement was beyond the observational capabilities of the time.

The purpose of the present paper is to examine the impact of the next generation of 30–60-m Extremely Large Telescopes (ELTs) on the feasibility of determining \dot{z} . The key issue is obviously the accuracy to which one can determine redshifts. In the absence of systematic instrumental effects this accuracy depends only on the intrinsic sharpness of the spectral features used, the number of features available and the signal-to-noise ratio (S/N) at which they are recorded (assuming that the features are resolved). In a photon-noise limited experiment the latter in turn depends only on the flux density of the source(s), the size of the telescope, the total combined telescope/instrument efficiency and the integration time. In this paper we will investigate this parameter space in detail,

expanding on previous work in several ways: in Section 3 we confirm the usefulness of the Ly α forest by quantifying the peculiar motions of the absorbing gas using hydrodynamic simulations of the intergalactic medium (IGM). In Section 4 we use Monte Carlo simulations of the Ly α forest to quantify how its properties translate to a radial velocity accuracy, and we consider the usefulness of other absorption lines in Section 5. These results are then used in Section 7 where we explore the observational parameter space and realistically assess the feasibility of a \dot{z} experiment with an ELT. Finally, we summarise our findings in Section 8.

2 MEASURING THE DYNAMICS

In any metric theory of gravity one is led to a very specific form of the metric by simply assuming that the Universe is homogeneous and isotropic. The evolution in time of this so-called Robertson-Walker metric is entirely specified by its global scale factor, $a(t)$. The goal is to measure or reconstruct this function. Recall that we observe the change of a with time by its wavelength-stretching effect on photons traversing the Universe. A photon emitted by some object at comoving distance χ at time t_{em} and observed by us at t_{obs} suffers a redshift of

$$1 + z(t_{\text{obs}}, t_{\text{em}}) = \frac{a(t_{\text{obs}})}{a(t_{\text{em}})} \quad (1)$$

(where, obviously, only two of the three variables χ , t_{em} and t_{obs} are independent). If it were possible to measure not only a photon's redshift but also its t_{em} , then the problem would be solved by simply mapping out the present-day relation between redshift and look-back time, i.e. $z(t_{\text{obs}} = t_0, t_{\text{em}})$, where t_0 denotes today.

A different approach to the problem is to consider how the redshift of an object at a fixed comoving distance χ evolves with t_{obs} , i.e. to consider the function $z_{|\chi}(t_{\text{obs}})$. For a given χ , t_{obs} determines t_{em} and so we have dropped the dependence on t_{em} . In principle, it is possible to map out $z_{|\chi}(t_{\text{obs}})$ (at least for $t_{\text{obs}} > t_0$), and this would be the most direct determination of $a(t)$, no matter which object is used. However, a full characterisation of $z_{|\chi}(t_{\text{obs}})$ would require observations over several Gyr. Over a much shorter timescale, Δt_{obs} , one can at most hope to determine the first order term of the Taylor expansion

$$\frac{dz_{|\chi}}{dt_{\text{obs}}}(t_{\text{obs}}) \approx \frac{z_{|\chi}(t_{\text{obs}} + \Delta t_{\text{obs}}) - z_{|\chi}(t_{\text{obs}})}{\Delta t_{\text{obs}}}. \quad (2)$$

As we will see presently, it turns out that measuring $dz_{|\chi}/dt_{\text{obs}}$ is in fact sufficient to reach our goal of reconstructing $a(t)$. By differentiating equation (1) with respect to t_{obs} we find

$$\frac{dz_{|\chi}}{dt_{\text{obs}}}(t_{\text{obs}}) = [1 + z_{|\chi}(t_{\text{obs}})]H(t_{\text{obs}}) - H(t_{\text{em}}), \quad (3)$$

where we have used that $dt_{\text{obs}} = [1 + z(t_{\text{obs}})]dt_{\text{em}}$ for a fixed χ , and $H = \dot{a}a^{-1}$. Evaluating at $t_{\text{obs}} = t_0$, replacing the unknown t_{em} with its corresponding redshift, and dropping the reminder that we are considering the redshift of an object at a fixed distance χ we simply obtain (McVittie 1962):

$$\dot{z} \equiv \frac{dz}{dt_{\text{obs}}}(t_0) = (1 + z)H_0 - H(z). \quad (4)$$

\dot{z} is a small, systematic drift as a function of time in the redshift of a cosmologically distant source as observed by us today. This effect is induced by the de- or acceleration of the expansion, i.e. by the change of the Hubble parameter H . Since H_0 is known




Figure 1. Redshift evolution of three objects with present-day redshifts $z(t_0) = 0.5, 3$ and 8 as a function of time of observation, for three different combinations of Ω_M and Ω_Λ as indicated. For each case the dotted lines indicate the Big Bang. t_{obs} is shown relative to the present day, t_0 .

(Freedman et al. 2001), this drift is a direct measure of the expansion velocity at redshift z . Measuring \dot{z} for a number of objects at different z hence gives us the function $\dot{a}(z)$. The point is that given $a(z)$ and $\dot{a}(z)$, one can reconstruct $a(t)$. A measurement of $\dot{z}(z)$ therefore amounts to a purely dynamical reconstruction of the expansion history of the Universe.

Predicting the redshift drift $\dot{z}(z)$ requires a theory of gravity. Inserting the Robertson-Walker metric into the theory's field equation results in the Friedman equation, which specifically links the expansion history with the densities, Ω_i , and equation of state parameters, w_i , of the various mass-energy components of the Universe. In the case of general relativity the Friedman equation is given by:

$$H(z) = H_0 \left[\sum_i \Omega_i (1+z)^{3(1+w_i)} + \Omega_k (1+z)^2 \right]^{\frac{1}{2}}, \quad (5)$$

where $\Omega_k = 1 - \sum_i \Omega_i$. Here we consider only two components: cold (dark) matter (CDM) with $w_M = 0$ and dark energy in the form of a cosmological constant with $w_\Lambda = -1$. In Fig. 1 we plot $z_{|\chi}(t_{\text{obs}})$ for three different objects [chosen to have $z(t_0) = 0.5, 3$ and 8] and for three different combinations of Ω_M and Ω_Λ , where we have also assumed $H_0 = 70 \text{ h}_{70} \text{ km s}^{-1} \text{ Mpc}^{-1}$. For each object, the redshift goes to infinity at some time in the past when the object first entered our particle horizon. If $\Omega_\Lambda = 0$ the redshift continually decreases thereafter as the expansion is progressively slowed down by Ω_M . Hence, in this case the redshifts of all objects are decreasing at the present time.

However, if $\Omega_\Lambda \neq 0$ the initial decrease is followed by a subsequent rise due to Λ relieving matter as the dominant mass-energy component and causing the expansion to accelerate. The turn-around point may lie either in the past or in the future, depending on the object's distance from us; i.e. if $\Omega_\Lambda \neq 0$ an object's redshift may be either increasing or decreasing at the present time. For distant (nearby) objects, the Universe was mostly matter (Λ)-dominated during the interval $[t_{\text{em}}, t_0]$ and hence underwent a net deceleration (acceleration), resulting in $\dot{z} < (>) 0$ at the present time (see Gudmundsson & Björnsson 2002 for a more detailed discussion).




Figure 2. The solid (dotted) lines and left (right) axis show the redshift drift \dot{z} (\dot{v}) as a function of redshift for various combinations of Ω_M and Ω_Λ as indicated. The dashed line shows \dot{z} for the case of dark energy having a constant $w_X = -2/3$ (and $\Omega_M, \Omega_X = 0.3, 0.7$).

These features are evident in Fig. 2, where we plot the expected present-day redshift drift, $\dot{z}(z)$, for various values of Ω_M and Ω_Λ (solid lines), and for a case where the dark energy $w \neq -1$ (dashed line). The redshift drift is also shown in velocity units (dotted lines), where $\dot{v} = c \dot{z} (1+z)^{-1}$. As noted above, the existence of a redshift region where $\dot{z} > 0$ is the hallmark of $\Omega_\Lambda \neq 0$. Note also the scale of Fig. 2. At $z = 4$ the redshift drift is of order 10^{-9} or 6 cm s^{-1} per decade. For comparison, the long-term accuracy achieved in extra-solar planet searches with the high-resolution echelle spectrograph HARPS on the ESO 3.6-m Telescope is of order 1 m s^{-1} (e.g. Lovis et al. 2005).

3 CHOOSING AN ACCELEROMETER

A priori, it is not at all obvious which spectral features of which set or class of objects might be best suited for a \dot{z} measurement. Clearly though, potential candidate targets should boast as many of the following desirable attributes as possible. (i) They should faithfully trace the Hubble flow. Although peculiar motions are expected to be random with respect to the Hubble flow, the additional noise introduced by them could potentially conceal the cosmic signal (Phillipps 1982; Teuber 1986; but see also Appendix A). (ii) The targets should have the sharpest possible spectral features to minimise the error on individual redshift measurements. (iii) The number of useful spectral features per target should be as high as possible in order to maximise the amount of relevant information per unit observing time. (iv) The targets should be as bright as possible and finally, (v) they should exist over a wide redshift range, and particularly at high z , where the signal is expected to be largest.

Clearly, several of these features are in conflict with each other. Requirements (i) and (ii) are conflicting because sharp spectral features require cold material which is generally found in dense regions inside deep potential wells, which in turn generate large peculiar accelerations. Similarly, point (i) clashes with point (iv) because high intrinsic luminosities require very massive and/or highly energetic processes, again implying deep potential wells. Since very dense regions are relatively rare there is likewise ten-

sion between requirements (ii) and (iii). Finally, the demands for brightness and high redshift are also difficult to meet simultaneously.

Hence, it seems that it is impossible to choose a class or set of objects that is the ideal target, in the sense that it is superior to every other class or set in each of the above categories. However, there is one class of ‘objects’ that meets all but one of the criteria.

3.1 The Ly α forest

The term ‘Ly α forest’ refers to the plethora of absorption lines observed in the spectra of all quasi-stellar objects (QSOs) shortwards of the Ly α emission line.¹ Almost all of this absorption arises in intervening intergalactic H I between us and the QSO (see Rauch 1998 for a review). Since the absorbing gas is physically unconnected with the background source against which it is observed we elegantly avoid the conflict between requirements (i) and (iv) above. However, as the gas is in photoionization equilibrium with the intergalactic ultraviolet (UV) background, its temperature is of order 10^4 K (Theuns, Schaye & Haehnelt 2000; Schaye et al. 2000). Consequently, the absorption lines are not particularly sharp and the typical line width is $\sim 30 \text{ km s}^{-1}$ (e.g. Kim, Cristiani & D’Odorico 2001). On the other hand, QSOs are among the brightest sources in the Universe and exist at all redshifts out to at least ~ 6 (Fan et al. 2006). Furthermore, each QSO spectrum at $z \gtrsim 2$ shows on the order of 10^2 absorption features.

In the following we will consider the question whether, apart from the Hubble expansion, other evolutionary processes acting on the absorbing gas might also significantly affect the measured positions of the absorption lines over the timescale of a decade or so. First and foremost is the issue to what extent the Ly α forest is subject to peculiar motions.

3.1.1 Peculiar motions

Cosmological hydrodynamic simulations (Cen et al. 1994; Zhang, Anninos & Norman 1995; Hernquist et al. 1996; Miralda-Escudé et al. 1996; Cen & Simcoe 1997; Charlton et al. 1997; Theuns et al. 1998; Zhang et al. 1998), analytic modelling (e.g. Bi & Davidsen 1997; Viel et al. 2002) and observations of the sizes and shapes of the absorbers (Bechtold et al. 1994; Dinshaw et al. 1994; Smette et al. 1995; Charlton, Churchill & Linder 1995; D’Odorico et al. 1998; Liske et al. 2000; Rollinde et al. 2003; Becker, Sargent & Rauch 2004; Coppolani et al. 2006; see also Rauch & Haehnelt 1995) all suggest that the Ly α forest absorption occurs in the large-scale, filamentary or sheet-like structures that form the cosmic web. These are at most mildly overdense and are participating, at least in an average sense, in the general Hubble expansion. However, at least some fraction of the absorbing structures (depending on redshift) must be expected to have broken away from the Hubble flow and to have begun collapsing under the influence of local gravitational potential wells.

This exact issue was investigated by Rauch et al. (2005) who studied the distribution of velocity shear between high-redshift Ly α absorption common to adjacent lines of sight separated by 1 to

¹ Since the Ly α forest is only observable from the ground for $z \gtrsim 1.7$ we only consider the high- z Ly α forest in this paper. We also exclude the higher column density Ly limit and damped Ly α (DLA) absorbers from the discussion.

300 h_{70}^{-1} kpc. They showed that the observed shear distributions were indeed in good agreement with the absorbing structures undergoing large-scale motions dominated by the Hubble flow. In fact, the distributions could be reproduced very convincingly by artificial pairs of spectra created from a hydrodynamical simulation of the IGM. This result instills us with further confidence that such simulations accurately capture the kinematics of the gas responsible for the Ly α forest.

Thus assured, we will now explicitly examine the peculiar velocities and accelerations of the absorbing gas in a hydrodynamic simulation of the IGM produced by the parallel TreeSPH code GADGET-2 (Springel 2005). We have used this code in its TreePM mode in order to speed up the calculation of long-range gravitational forces. The simulations were performed with periodic boundary conditions and with 400^3 dark matter and 400^3 gas particles in a box of 60 h_{100}^{-1} Mpc comoving size. Radiative cooling and heating processes were followed using an implementation similar to that of Katz, Weinberg & Hernquist (1996) for a primordial mix of hydrogen and helium. The UV background was taken from Haardt & Madau (1996). In order to further increase the speed of the simulation we applied a simplified star formation criterion: all of the gas at overdensities $> 10^3$ times the mean density and with temperature $< 10^5$ K was turned into stars. The cosmological parameters were set to $\Omega_M = 0.26$, $\Omega_\Lambda = 0.74$, $\Omega_b = 0.0463$, $n_s = 0.95$, $\sigma_8 = 0.85$ and $H_0 = 72 \text{ km s}^{-1} \text{ Mpc}^{-1}$, where Ω_b is the baryonic density parameter, and n_s and σ_8 are the spectral index and amplitude of the linear dark matter power spectrum, respectively. These values are in excellent agreement with recent joint analyses of CMB, SNIa, galaxy clustering and Ly α forest data (Spergel et al. 2007; Seljak et al. 2006; Viel, Haehnelt & Lewis 2006). The Λ CDM transfer function was computed with CMBFAST (Seljak & Zaldarriaga 1996). The above corresponds to the simulation series B2 of Viel, Haehnelt & Springel (2004) which has been widely used for cosmological studies. Moreover, this is the exact same simulation that successfully reproduced the velocity shear distributions observed by Rauch et al. (2005).

We pierced our simulation box with 1000 random lines of sight (LOS) and noted the physical properties of the absorbing gas along these lines at three different redshifts ($z = 2, 3$ and 4). In Figs. 3 and 4 we plot the distributions of peculiar velocity and acceleration, respectively. The acceleration of the gas is computed from its velocity by dividing the latter by a dynamical time $t_{\text{dyn}} = (G\rho)^{-1/2}$ (e.g. Schaye 2001), where ρ is the total matter density. Evidently, the distributions do not evolve rapidly with redshift. Since a given parcel of gas cannot be decelerated, peculiar velocities cannot decrease with time and hence their distribution shifts to slightly larger values at lower redshifts. The average peculiar acceleration, on the other hand, decreases with time because by volume the Universe is dominated by low-density regions in which the gas density decreases as the Universe expands.

Peculiar motions are expected to be randomly oriented with respect to the LOS. Hence, when averaging over a large number of individual z measurements, peculiar motions will only introduce an additional random noise component but no systematic bias. In Appendix A we explicitly derive an expression for the observed redshift drift in the presence of peculiar motions. This expression can be used to translate the distributions in Figs. 3 and 4 into the corresponding error distribution on \dot{z} . We find that for a decade-long experiment the error due to peculiar motion is of order $\sim 10^{-3} \text{ cm s}^{-1}$. This must be compared to the error induced by photon noise for an individual \dot{z} measurement from a *single* absorption line. Clearly, if we wish to be able to detect the redshift drift




Figure 3. Probability distribution functions (PDFs) of the peculiar velocity of the absorbing gas along 1000 random lines of sight through our simulation box at three different redshifts as indicated. Note that we are *not* using the value of the velocity's component along the LOS, but rather the modulus of the full 3-dimensional velocity.




Figure 4. As Fig. 3 for the peculiar acceleration.

over a decade or so the overall accuracy of the whole experiment has to be of order $\sim 1 \text{ cm s}^{-1}$. Since this will be achieved using hundreds of absorption lines the error on an individual \dot{z} measurement will be at least a factor of ~ 10 larger. The error due to peculiar motions is therefore sub-dominant by at least ~ 4 orders of magnitude. Hence we conclude that peculiar motions will have no detrimental impact whatsoever on a redshift drift experiment targeting the Ly α forest.

3.1.2 Galactic feedback

So far we have only considered peculiar motions induced by gravity. In principle, there are of course a number of non-gravitational ways of imparting kinetic energy to the absorbing gas, mostly involving galactic feedback. The main reason to believe that galactic feedback must have had a far-reaching impact on the Ly α forest is its early and widespread low-level metal enrichment, even at fairly low column densities (e.g. Cowie et al. 1995; Ellison et al. 2000; Schaye et al. 2000, 2003; Aguirre et al. 2004; Simcoe, Sargent & Rauch 2004). Although very low-level enrichment could be achieved in situ with Population III stars (e.g. Gnedin & Ostriker 1997) the general consensus is that some sort of mechanism is required to transport metals from (proto-)galaxies into the IGM. In contrast, no consensus has been reached as to which of the possible mechanisms (or combination thereof) is the correct one. Candidates are mergers and tidal interactions (e.g. Gnedin 1998), ram pressure stripping, radiation pressure on dust grains (e.g. Aguirre et al. 2001), photo-evaporation during reionization (Barkana & Loeb 1999) and SN-driven winds, either from low-mass, (pre-)galactic halos at $z \approx 10$ (e.g. Madau et al. 2001) or from massive starbursting galaxies at $z < 5$ (e.g. Aguirre et al. 2001; Adelberger et al. 2003). From our point of view the last possibility is the most worrying as it has the highest potential of significantly altering the kinematic structure of the IGM at the time of observation.

There exists persuasive evidence of the existence of galactic superwinds from Lyman break galaxies at $z \approx 3$ (e.g. Pettini et al. 2001) and it is likely that some fraction of strong metal absorption lines are connected with these structures (e.g. Adelberger et al. 2003; Simcoe et al. 2006). On the other hand, there is no observational evidence at all that superwinds are significantly stirring up the high redshift Ly α forest at the time we observe it. Using column density and optical depth differences across a close pair of lines of sight Rauch et al. (2001) found no indication in the forest's small-scale density structure for widespread recent disturbances. Similarly, as discussed above, the velocity shear between adjacent lines of sight is entirely explained by the Hubble flow and gravitational instability (Rauch et al. 2005), leaving little room for non-gravitationally induced motion. Indeed, the simple fact that the aforementioned hydrodynamical simulations – which did *not* include any galactic feedback – were so successful in reproducing the observed properties of the Ly α forest, including its line broadening distribution and clustering, raises the question of how significant amounts of feedback could be integrated without upsetting the existing agreement between the models and the data (Theuns, Mo & Schaye 2001). It seems that the volume filling factor of galactic superwinds is limited to a few per cent (Theuns et al. 2002; Desjacques et al. 2004; Pieri & Haehnelt 2004; Cen et al. 2005; Bertone et al. 2005). We conclude that, whatever the process of metal enrichment may be, there is currently no reason to believe that it has a wholesale effect on the kinematics of the general IGM (as probed by the Ly α forest) at the time of observation.

3.1.3 Optical depth variations

Consider the gas responsible for a given Ly α forest absorption feature. If the physical properties of the gas change over the timescale of a decade or so this will cause a variation of the feature's optical depth profile. Potentially, this could in turn lead to a small shift in the feature's measured position and hence mimic a redshift drift. The precise magnitude of this additional \dot{z} error component will

depend on the method used to extract the signal but we can gain an impression of the relevance of the effect by comparing the expected optical depth variation to the apparent optical depth change (at a fixed spectral position) induced by the redshift drift.

The gas properties we consider here are density, temperature and ionization fraction. In the fluctuating Gunn Peterson approximation (e.g. Hui, Gnedin & Zhang 1997) the Ly α optical depth is related to these quantities by $\tau \propto (1 + \delta)^2 T^{-0.7} \Gamma^{-1}$, where δ and T are the gas overdensity and temperature, respectively, and Γ is the photoionization rate. The interplay between photoionization heating of the gas and adiabatic cooling leads to a tight relation between temperature and density, which can be well approximated by $T = T_0(1 + \delta)^\gamma$ (Hui & Gnedin 1997). Hence we obtain

$$\tau \propto (1 + \delta)^{2-0.7\gamma} T_0^{-0.7} \Gamma^{-1}. \quad (6)$$

How do these quantities evolve with time? According to linear theory density perturbations grow as $(1 + z)^{-1}$. Although exactly true only for an Einstein-de Sitter Universe, this represents an upper limit in open and flat models with a cosmological constant. Hence we will err on the side of caution by adopting this growth factor in the following. The evolution of the temperature-density relation can be gleaned from figure 6 of Schaye et al. (2000): $d\gamma/dz = -0.1$ and $dT_0/dz = 0.25 T_0$ (where we have conservatively used the steep evolution between redshifts 2 and 3). We take the evolution of Γ from figure 7 of Bolton et al. (2005): $d\Gamma/dz = -0.3\Gamma$. Taking the appropriate derivatives of equation (6), inserting the above values and translating a decade in our reference frame to a redshift difference dz at $z = 3$ we find an optical depth variation of $d\tau = 3 \times 10^{-9} f \tau$, where $f \approx 0.2, 0.2, 0.05, 0.3$ for the variation of δ, T_0, γ and Γ , respectively. As we will see in Section 4.4, these values for $d\tau$ are at least ~ 2 orders of magnitude smaller than the typical optical depth changes due to the redshift drift. We therefore conclude that changes in the physical properties of the absorbing gas are not expected to interfere with a \dot{z} measurement from the Ly α forest.

3.2 Molecular absorption lines

Before we move on to investigate the details of a redshift drift experiment using ELT observations of the Ly α forest, let us briefly digress here to consider a very different \dot{z} experiment using another future facility. We have explored in some detail the possibility of using the Atacama Large Millimeter Array (ALMA) to measure \dot{z} from rotational molecular transitions seen in absorption against background continuum sources. Going to the (sub-)mm regime has the advantages of potentially very high resolution and less photon noise for a given energy flux. Furthermore, molecular absorption lines can be very sharp, with line widths as low as $\lesssim 1 \text{ km s}^{-1}$. However, the molecular gas in nearby galaxies is strongly concentrated towards the central regions. Hence we must expect the gas to be subject to peculiar accelerations similar in magnitude to the cosmological signal. That in itself would not necessarily be problematic as long as we had many individual \dot{z} measurements from different objects over which to average. Unfortunately, it seems unlikely that this will be the case. At present, rotational molecular lines have been detected in only four absorption systems (with redshifts 0.25–0.89; Wiklind & Combes 1999, and references therein), despite intensive searches (e.g. Curran et al. 2004). Based on the incidence of these lines and the number of continuum sources with fluxes larger than 10 mJy at 90 GHz we estimate that the number of molecular absorption systems observable with ALMA will be ~ 50 , with only 5–10 of these showing narrow lines – not enough to overcome the

uncertainties due to peculiar motions. Hence we have decided not to pursue the case for molecular absorption lines any further.

4 SENSITIVITY OF THE $\text{Ly}\alpha$ FOREST TO RADIAL VELOCITY SHIFTS

In Section 2 we have seen that the redshift drift is a very small effect. In order to detect it, an experiment must achieve an overall accuracy with which radial velocity shifts can be determined of order $\sim 1 \text{ cm s}^{-1}$. In this section we will investigate how the properties of the $\text{Ly}\alpha$ forest translate to a radial velocity accuracy, σ_v , and how σ_v depends on the instrumental characteristics of the spectra and on redshift. Specifically, we would like to know how many $\text{Ly}\alpha$ forest spectra of which resolution and S/N are needed at a given redshift to achieve the required σ_v .

We will investigate these issues using artificial spectra. High-resolution observations have demonstrated that, to first approximation, the $\text{Ly}\alpha$ forest can be decomposed into a collection of individual absorption lines (e.g. Kim et al. 2001). These are usually taken to be Voigt profiles and so each line is characterised by three parameters: redshift, z , H I column density, N_{HI} , and velocity width, b . Here, we will reverse this decomposition process and generate (normalised) spectra with the desired instrumental characteristics from given lists of absorption lines. We will use two types of line lists. First, we will generate line lists from Monte Carlo (MC) simulations based on the statistics of the largest available samples of absorption lines. Secondly, to validate our simulations, we will use 8 line lists available in the literature that have previously been derived from high-resolution observations.

4.1 Simulated absorption line lists

We form simulated MC line lists by simply randomly drawing values for the absorption line parameters from their observed distributions (e.g. Hu et al. 1995; Lu et al. 1996; Kirkman & Tytler 1997; Kim et al. 1997, 2001, 2002):

$$f(z, N_{\text{HI}}, b) \propto (1+z)^\gamma N_{\text{HI}}^{-\beta} \exp\left[-\frac{(b-\bar{b})^2}{2\sigma_b^2}\right], \quad (7)$$

where $\gamma = 2.2$, $\beta = 1.5$, $\bar{b} = 30 \text{ km s}^{-1}$ and $\sigma_b = 8 \text{ km s}^{-1}$. We impose limits of $15 < b < 100 \text{ km s}^{-1}$ and also restrict N_{HI} to the classical $\text{Ly}\alpha$ forest regime, excluding Lyman limit and DLA systems: $12 < \log N_{\text{HI}}(\text{cm}^{-2}) < 16$. The above distribution is normalised to give 10^2 lines with $13.64 < \log N_{\text{HI}}(\text{cm}^{-2}) < 16$ per unit redshift at $z = 2$ (Kim et al. 2001). The actual number of absorption lines in a given line list is drawn from a Poisson distribution with a mean determined by the normalisation.

The above MC approach allows us to quickly generate large amounts of spectra with realistic characteristics. Note, however, that we make no assumptions regarding the underlying physics of the IGM in which the absorption occurs. We simply use the observational fact that $\text{Ly}\alpha$ forest spectra can be well represented as a random collection of Voigt profiles.

The most significant difference between our MC line lists and the real $\text{Ly}\alpha$ forest is clustering: the real $\text{Ly}\alpha$ forest is not randomly distributed in redshift but shows significant redshift-space correlations on scales of at least 100 km s^{-1} (e.g. Cristiani et al. 1995; Fernández-Soto et al. 1996; Liske et al. 2000). The impact of clustering will be discussed in detail in Section 4.6.

Table 1. Observed $\text{Ly}\alpha$ forest line lists from the literature.

QSO	z_{QSO}	λ range ^a	$N_{\text{Ly}\alpha}^b$	Reference
Q1101–264	2.145	3226–3810	290	1
J2233–606	2.238	3400–3850	226	2
HE1122–1648	2.400	3500–4091	354	1
HE2217–2818	2.413	3550–4050	262	2
HE1347–2457	2.617	3760–4335	362	1
Q0302–003	3.281	4808–5150	223	1
Q0055–269	3.655	4852–5598	535	1
Q0000–26	4.127	5380–6242	431	3

References: 1 = Kim et al. (2002), 2 = Kim et al. (2001), 3 = Lu et al. (1996).

^aWavelength range covered by the $\text{Ly}\alpha$ forest line lists in Å.

^bNumber of $\text{Ly}\alpha$ forest lines.

4.2 Real absorption line lists

We have collected 8 QSO absorption line lists from the literature (see Table 1). These were derived from UVES/VLT (7 objects) or HIRES/Keck data (1 object). All spectra have a resolution of $\text{FWHM} \approx 7 \text{ km s}^{-1}$, while the typical S/N per pixel varies from ~ 10 for Q0000–26 to ~ 50 in the case of HE2217–2818. In all cases the absorption line lists were derived by fitting Voigt profiles to the spectra using VPFIT.² Details of the data acquisition and reduction, as well as of the line fitting and identification processes are given by Kim et al. (2001, 2002) and Lu et al. (1996).

We note that two of the spectra, those of Q1101–264 and Q0000–26, contain DLAs, which effectively block out parts of the spectra. However, this is only expected to have a significant effect at high redshift, where the $\text{Ly}\alpha$ forest line density is high, and so we have excluded the affected spectral region only in the case of Q0000–26.

4.3 The second epoch

Simulating a z measurement requires a second epoch ‘observation’ of the same line of sight: we generate a second epoch absorption line list from the original line list (both real and simulated) by simply shifting the redshift of each line according to a given cosmological model:

$$\Delta z_i = \dot{z}(z_i; H_0, \Omega_M, \Omega_\Lambda) \Delta t_0, \quad (8)$$

where Δt_0 is the assumed time interval between the first and second epochs. Unless stated otherwise, we use the standard, general relativistic cosmological model, assuming fiducial parameter values of $H_0 = 70 \text{ h}_{70} \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_M = 0.3$ and $\Omega_\Lambda = 0.7$.

4.4 Generating spectra

Given an absorption line list a normalised spectrum is generated by

$$S(\lambda) = \exp\left\{-\sum_i^{N_{\text{al}}} \tau[\lambda_\alpha(1+z_i), N_{\text{HI},i}, b_i]\right\}, \quad (9)$$

where λ is the observed wavelength, N_{al} is the number of absorption lines in the spectrum, τ is the optical depth of a Voigt profile and $\lambda_\alpha = 1215.67 \text{ Å}$ is the rest wavelength of the H I $\text{Ly}\alpha$

² by R.F. Carswell et al., see <http://www.ast.cam.ac.uk/~rfc/vpfit.html>.




Figure 5. (a) Example of an artificial Ly α forest spectrum at $z \approx 3$ generated from a Monte Carlo absorption line list. (b) Close-up of the region around 5000 Å. (c) Artificial spectrum generated from the observed line list of Q0302–003. Both spectra have S/N = 100.

transition. This spectrum is then pixelised, using a pixel size of 0.0125 Å, and convolved with a Gaussian line-spread function. Unless stated otherwise, we will use a resolution element four times the pixel size, corresponding to a resolution of $R = 100\,000$ at 5000 Å. We then add random noise to the spectrum assuming Poisson statistics, i.e. we assume a purely photon-noise limited experiment. All S/N values quoted in this paper refer to the S/N per pixel in the continuum. To begin with, we will only consider the spectral range between the assumed background QSO’s Ly α and Ly β emission lines, i.e. the classical Ly α forest region. By construction this region cannot contain any H I transitions that are of higher order than Ly α , and it covers an absorption redshift range of $\lambda_\beta/\lambda_\alpha(1 + z_{\text{QSO}}) - 1 < z < z_{\text{QSO}}$, where z_{QSO} is the QSO’s redshift and $\lambda_\beta = 1025.72$ Å is the rest wavelength of the H I Ly β transition. In Sections 5.1 and 5.2 we will also consider other spectral regions.

In Fig. 5 we show examples of artificial spectra generated from both simulated (panels a and b) and real line lists (panel c). The expected flux difference between two spectra of the same absorption lines, taken a decade apart, is shown in Fig. 6. A similar plot for the optical depth difference reveals that the variations discussed in Section 3.1.3 are unproblematic.

4.5 Defining σ_v

In principle, a z determination will involve the measurement of radial velocity differences between the corresponding features of a pair of spectra of the same object taken several years apart. We now need a method to estimate the accuracy to which these differences can be determined. This requires knowledge of how exactly the measurement will be performed. However, a priori it is not at all obvious what the optimal signal extraction method might be. To proceed nevertheless we choose to base our analysis on the generic concept of the total radial velocity information content of a spectrum, which was developed by Bouchy, Pepe & Queloz (2001) in




Figure 6. Flux difference between two artificial, noiseless spectra of the same Ly α forest at $z \approx 3$ simulated for two observing epochs separated by $\Delta t_0 = 10$ yr and for various combinations of Ω_M and Ω_Λ as indicated. The redshift drift implied by these parameters is also given.

the context of optimising radial velocity searches for extra-solar planets.

Following Bouchy et al. (2001) we begin by expressing the flux observed in pixel i at the second epoch as a small perturbation on the first epoch flux in the same pixel:

$$S_{2i} = S_{1i} + \frac{dS_i}{d\lambda} \frac{\Delta v_i}{c} \lambda_i, \quad (10)$$

which defines a small velocity shift Δv_i for each pixel. λ_i is the observed wavelength of the i th pixel and $dS_i/d\lambda$ is the spectral slope of the flux at that pixel. To first order the slope does not change between the two epochs and hence it carries no epoch designation. Averaging the velocity shift over all pixels in a spectrum, using weights w_i , we have:

$$\Delta v = \frac{\sum_i \Delta v_i w_i}{\sum_i w_i}. \quad (11)$$

Clearly, the weight for the i th pixel should be chosen as the inverse variance of Δv_i . In calculating this variance we must differ from Bouchy et al. (2001). In the case of stars one of the spectra can be assumed to be a perfect, noiseless template, essentially because additional information on the same type of star can be used to define it. However, since every Ly α forest spectrum is unique we cannot make the same assumption here, so that in our case both spectra have noise. Hence we find:

$$\sigma_{v_i}^2 = \left(\frac{c}{\lambda_i \frac{dS_i}{d\lambda}} \right)^2 \left[\sigma_{1i}^2 + \sigma_{2i}^2 + \frac{(S_{2i} - S_{1i})^2}{\left(\frac{dS_i}{d\lambda} \right)^2} \sigma_{S'_i}^2 \right], \quad (12)$$

where σ_{1i} and σ_{2i} are the flux errors in the i th pixel of the first and second epoch spectra, respectively, and $\sigma_{S'_i}$ is the error on the slope of the flux at pixel i . We can see that a low weight is assigned to noisy pixels and to those that have a small gradient, i.e. pixels in the continuum or in the troughs of saturated absorption lines. Finally, with the above choice of weights the error on Δv is given by:

$$\sigma_v^2 = \frac{\sum_i \sigma_{v_i}^2 w_i^2}{\left(\sum_i w_i \right)^2} = \frac{1}{\sum_i \sigma_{v_i}^{-2}}. \quad (13)$$

The above process has the advantage of conveniently attaching a single figure of merit to a given pair of spectra in a non-parametric, model-independent way: σ_v simply represents the fun-

damental photon-noise limit of the accuracy to which an overall velocity shift between the two spectra can be determined. It is essentially just a measure of the ‘wiggleness’ and of the S/N of the spectra. However, we point out that σ_v does not entirely capture all of the information contained in a pair of spectra with respect to a z measurement. From Fig. 2 it is clear that the difference between the first and second epoch spectra is not simply an overall velocity shift. The second epoch spectrum will also be slightly compressed with respect to the first epoch spectrum because the redshift drift is larger at higher redshifts (i.e. longer wavelengths) than at lower redshifts. The additional information that is contained in this small alteration of the spectrum’s shape is entirely ignored by the above method [because of the simple averaging operation in equation (11)] and hence it is clearly a sub-optimal method of estimating the sensitivity of a pair of spectra to \dot{z} . Despite this shortcoming we will retain the above definition of σ_v for the sake of its simplicity.

4.6 Results

We are now ready to derive the relevant scaling relations for σ_v . We begin by making two points regarding equation (12). First, we note that for a fixed total integration time (the sum of the integration times spent observing the first and second epoch spectra) the sum $\sigma_{1i}^2 + \sigma_{2i}^2$ in equation (12) takes on its minimum when the first and second epoch integration times are equal. Hence, the smallest possible σ_v is only achieved when the spectra of both epochs have the same S/N. In the following we will assume that this is the case. Secondly, from equation (12) it is clear that σ_v scales as $(\text{S/N})^{-1}$, as expected for a photon-noise limited experiment.

Consider now a set of N_{QSO} targets that all lie at the same redshift, each of which has been observed at two epochs such that all $2N_{\text{QSO}}$ spectra have the same S/N. Again, since we are considering a purely photon-noise limited experiment, σ_v should scale as $N_{\text{pix}}^{-1/2}$ (where N_{pix} is the total number of independent data points in the sample) and hence also as $N_{\text{QSO}}^{-1/2}$. Furthermore, in this ideal case, it is irrelevant how the *total* S/N is divided among the targets. Hence, for simplicity we will use $N_{\text{QSO}} = 1$ in the following.

We now examine the behaviour of σ_v as a function of redshift using the MC absorption line lists. For various QSO redshifts in the range $2 \leq z_{\text{QSO}} \leq 5$ we have generated 10 pairs of Ly α forest line lists and spectra as described in Sections 4.1 to 4.4, where each spectrum was given a S/N of 13 000. We then measured each pair’s σ_v according to equation (13). The result is shown as blue dots in Fig. 7, where each point and error bar is the mean and ± 1 r.m.s. of the 10 individual measurements at each redshift. We stress that we are plotting the expected accuracy of a velocity shift measurement performed on a single pair of spectra of a single target at a given redshift, where each spectrum has S/N = 13 000. We are *not* plotting the *combined* accuracy of 10 such pairs.

From Fig. 7 we can see that the radial velocity sensitivity improves rather rapidly with redshift for $z_{\text{QSO}} < 4$, but the decrease is somewhat shallower at $z_{\text{QSO}} > 4$. Overall, σ_v improves by a factor of almost 3 when moving from $z_{\text{QSO}} = 2$ to 5. Specifically, we find:

$$\sigma_v \propto \begin{cases} (1 + z_{\text{QSO}})^{-1.7} & z_{\text{QSO}} < 4 \\ (1 + z_{\text{QSO}})^{-0.9} & z_{\text{QSO}} > 4 \end{cases} \quad (14)$$

The behaviour of σ_v as a function of redshift is due to the combination of several factors. The first is the redshift evolution of the Ly α forest line density (cf. equation 7). At higher redshift more spec-




Figure 7. The blue dots with error bars show the accuracy with which a radial velocity shift can be determined from a Ly α forest spectrum as a function of QSO redshift. The red line parameterises the redshift dependence as in equation (14). Each point is the mean σ_v measured from 10 pairs of artificial spectra with S/N = 13 000, generated from simulated MC absorption line lists. The error bars show the ± 1 r.m.s. range of the 10 simulations. The green triangles show the results for simulations where the redshift evolution of the Ly α forest has been switched off, i.e. where $\gamma = 0$ (cf. equation 7). The orange crosses show the result of also restricting the σ_v measurement in each spectrum to a redshift path of constant length $\Delta z = 0.4$, as opposed to using the full Ly α forest region between the QSO’s Ly α and Ly β emission lines.

tral features are available for determining a velocity shift and so σ_v decreases. However, from $z \approx 4$ the absorption lines severely blanket each other and the number of sharp spectral features does not increase as rapidly anymore, causing the flattening of σ_v at $z_{\text{QSO}} > 4$. The green triangles in Fig. 7 show the σ_v measurements that result from simulations where the redshift evolution of the Ly α forest has been switched off, i.e. where the evolutionary index γ has been set to 0 (cf. equation 7). Indeed, we can see that in this case there is no evidence of a break.

Secondly, we recall that each Ly α forest spectrum covers the entire region between the QSO’s Ly α and Ly β emission lines. The redshift path length of this region is given by $\Delta z = 0.156(1 + z_{\text{QSO}})$. Hence the number of independent pixels per spectrum also increases as $(1 + z_{\text{QSO}})$. Since the S/N per pixel is kept constant this implies a larger number of photons per spectrum and hence an improved sensitivity to radial velocity shifts. The effect of this can be seen by comparing the green triangles in Fig. 7 with the orange crosses which are the result of using a constant redshift path length of $\Delta z = 0.4$ for each σ_v measurement, as well as $\gamma = 0$.

Finally, with $\gamma = 0$ and $\Delta z = \text{const}$, the sensitivity to *wavelength* shifts should be constant as a function of z_{QSO} , and so the sensitivity to *velocity* shifts should go as $(1 + z_{\text{QSO}})^{-1}$ (as can be seen from equation 12). In fact, the orange crosses in Fig. 7 decrease more slowly than this because the widths of the absorption lines in wavelength space increase as $(1 + z_{\text{QSO}})$, making the edges of the lines less steep and hence slightly decreasing the spectrum’s sensitivity to wavelength shifts.

In Fig. 8 we compare these results derived from the MC line lists to those from the real line lists. From Table 1 we can see that the real line lists (and hence the corresponding spectra) do not cover the full Ly α forest regions, with differently sized pieces missing both at the low and high redshift ends. Therefore we must correct the σ_v values derived from the real line lists in order to make them




Figure 8. Comparison of σ_v measurements derived from simulated MC line lists (blue dots and solid line, same as in Fig. 7) and from real line lists (red stars). The green circles show the results derived from simulated MC line lists that include a simple scheme for clustering absorption lines in redshift space (see text for details).

directly comparable to the values from the simulated lists. The correction is achieved by first assigning a new, slightly different QSO redshift to each spectrum, such that the ‘missing’ low and high redshift parts of the Ly α forest region are equally large. We then decrease the measured σ_v by a factor $(\Delta z_{\text{obs}}/\Delta z)^{1/2}$, where Δz_{obs} is the redshift path length covered by the observed line list and Δz is the redshift path length of the full Ly α forest region at the new z_{QSO} . The correction factors range from 0.56 to 0.99.

The red stars in Fig. 8 show the corrected σ_v values derived from single pairs of spectra generated from the real absorption line lists with $S/N = 13\,000$, while the blue dots show the measurements from the MC line lists (same as in Fig. 7). Overall the agreement between the results from the MC and real line lists is very good, particularly at high redshift. At $z_{\text{QSO}} \approx 2.4$ the σ_v values from the real line lists are ~ 15 per cent higher than those from the MC lists. By far the most significant deviation occurs at the lowest redshift where the σ_v of Q1101–264 is higher than expected by 47 per cent. However, this is not too surprising as the line of sight towards Q1101–264 is known to pass through an unusually low number of absorbers with $\log N_{\text{HI}}(\text{cm}^{-2}) > 14$ (Kim et al. 2002).

We believe that the small differences at $z_{\text{QSO}} \approx 2.4$ are mainly due to clustering of real absorption lines in redshift space. Clustering has the effect of reducing the number of spectral features because it increases line blanketing. However, at high redshift line blanketing is already severe because of the high line density and so clustering has a relatively smaller effect at high redshift than at low redshift. We demonstrate that clustering can explain the observed differences by generating a new set of MC line lists which incorporate a toy clustering scheme: first, we randomly draw the positions of ‘cluster’ centres from the Ly α forest redshift distribution. We then populate each ‘cluster’ with $n \geq 0$ absorbers, where n is drawn from a Borel distribution (Saslaw 1989). Since n can be 0, 1 or > 1 this process generates voids, single ‘field’ absorbers, as well as groups and clusters of lines. Finally, absorption lines are distributed around their ‘host’ clusters according to a Gaussian distribution with $\text{FWHM} = 120 \text{ km s}^{-1}$. The σ_v values that result from this new set of MC line lists are shown as open green circles in Fig. 8. The increase of σ_v compared to the unclustered simulations is clearly very similar to that observed for the real line lists




Figure 9. Radial velocity accuracy as a function of spectral resolution. Each point and error bar is the mean and ± 1 r.m.s. of 10 simulations at $z_{\text{QSO}} = 4$ and $S/N = 13\,000$. The pixel size is kept constant and is chosen such that a resolution element is sampled by 3 pixels at the highest resolution considered. The labels along the bottom axis denote the resolving power at 5000 Å , while the labels along the top axis show the equivalent FWHM of an unresolved line. The vertical red line marks the mode of the absorption lines’ b parameter distribution (cf. equation 7).

and we conclude that clustering can indeed explain the small difference between the results obtained from the observed and simulated line lists at $z_{\text{QSO}} \approx 2.4$. In any case, the near coincidence of the σ_v value of J2233–606 at $z_{\text{QSO}} = 2.23$ with the expected value demonstrates that not all lines of sight are adversely affected by clustering. In the following we will assume that such sightlines can be pre-selected and hence we will ignore the effects of clustering in the rest of this section.

We now turn to the behaviour of σ_v as a function of spectral resolution. For various resolving powers in the range $2500 \leq R \leq 1.33 \times 10^5$ we have generated 10 pairs of line lists and spectra with $z_{\text{QSO}} = 4$ and $S/N = 13\,000$, and measured their σ_v values as before. The result is presented in Fig. 9, where we show the resolving power along the bottom axis and the corresponding FWHM of an unresolved line along the top axis. We stress that the pixel size was the same for all spectra ($= 0.0125 \text{ Å}$) and that it was chosen such that a resolution element was well sampled even at the highest resolution considered. Hence the strong increase of σ_v towards lower R in Fig. 9 is not due to different numbers of photons per spectrum or sampling issues. Instead, it is simply due to the loss of information caused by convolving the spectrum with a line-spread function that is broader than the typical intrinsic absorption line width (marked by the vertical line in Fig. 9). Indeed, at $R \gtrsim 30\,000$ the Ly α forest is fully resolved and in this regime σ_v is independent of R .

Summarising all of the above we find that the accuracy with which a radial velocity shift can be determined from the Ly α forest scales as:

$$\sigma_v = 2 \left(\frac{S/N}{2370} \right)^{-1} \left(\frac{N_{\text{QSO}}}{30} \right)^{-\frac{1}{2}} \left(\frac{1+z_{\text{QSO}}}{5} \right)^{-1.7} \text{ cm s}^{-1}, \quad (15)$$

where the last exponent changes to -0.9 at $z_{\text{QSO}} > 4$, and where the same S/N (per pixel) is assumed for all N_{QSO} spectra at both epochs.

5 SENSITIVITY GAINS FROM OTHER SPECTRAL REGIONS

In the previous section we have investigated the sensitivity of the $\text{Ly}\alpha$ forest to radial velocity shifts, using only the H I $\text{Ly}\alpha$ transition in the region between a QSO's $\text{Ly}\alpha$ and $\text{Ly}\beta$ emission lines. However, modern echelle spectrographs are capable of covering a much wider spectral range in a single exposure and so the question arises whether other spectral regions, containing absorption lines from other ions or other H I transitions, can expediently contribute towards a \dot{z} measurement.

Before applying the procedure of the previous section to more extended spectra, we briefly follow-up on the discussion at the end of Section 4.5, where we pointed out that σ_v does not capture the information that is contained in the \dot{z} -induced change of the shape of a spectrum. In the case of a $\text{Ly}\alpha$ -only spectrum this change of shape consisted of a compression of the spectrum. If we allow additional transitions with different rest wavelengths then we no longer have a one-to-one correspondence between absorption redshift and wavelength, and so the redshift drift will in general induce a much more complex change of the shape of a spectrum. Hence it is clear that any attempt to harness this additional information must involve the complete identification and modelling of all absorption features used in the analysis.

Indeed, the complete identification of all metal absorption lines will be necessary in any case, even if one endeavours to measure the redshift drift only from the $\text{Ly}\alpha$ forest. The point is that the $\text{Ly}\alpha$ forest may of course be ‘contaminated’ by metal lines. Since these lines may arise in absorption systems that lie at completely different redshifts compared to the $\text{Ly}\alpha$ lines, their redshift drift may also be very different. Hence, any lines wrongly assumed to be $\text{Ly}\alpha$ could potentially lead to erroneous results. The only way to avoid such biases is to completely identify all absorption features.

5.1 The $\text{Ly}\beta$ forest

Each H I absorption line in the $\text{Ly}\alpha$ forest has corresponding counterparts at shorter wavelengths that result from the higher order transitions of the Lyman series. Obviously, all the arguments concerning the suitability of the $\text{Ly}\alpha$ forest for a \dot{z} measurement also apply to these higher order lines, and so, in principle, they should be almost as useful as the $\text{Ly}\alpha$ lines (where the qualifier ‘almost’ is owed to the decreasing optical depth with increasing order). However, in practice one has to contend with (i) confusion due to the overlap of low-redshift, low-order lines with high-redshift, high-order lines, and (ii) with an increased uncertainty in the placement of the QSO's continuum. Both of these difficulties are aggravated as one proceeds up the series towards shorter wavelengths. For these reasons we will not consider any H I transitions that are of higher order than $\text{Ly}\beta$.

We now extend our simulated spectra by adding the region between a QSO's $\text{Ly}\beta$ and $\text{Ly}\gamma$ emission lines immediately bluewards of the $\text{Ly}\alpha$ forest. This region extends the redshift path for $\text{Ly}\alpha$ lines by a factor of 1.28 towards lower redshifts and also contains $\text{Ly}\beta$ lines in the redshift range $\lambda_\gamma/\lambda_\beta(1+z_{\text{QSO}})-1 < z < z_{\text{QSO}}$, where $\lambda_\gamma = 972.54 \text{ \AA}$. Note, however, that the real $\text{Ly}\alpha$ line lists do not extend below the QSO's $\text{Ly}\beta$ emission lines (cf. Table 1). Hence we are forced to supplement the real $\text{Ly}\beta$ absorption lines in this region with $\text{Ly}\alpha$ lines drawn from the MC simulations.

In Fig. 10 we compare the σ_v measurements derived from the extended MC line lists (blue dots) to the $\text{Ly}\alpha$ -only results (solid line, same as in Figs. 7 and 8). We find that including the $\text{Ly}\beta$ for-




Figure 10. Comparison of σ_v measurements derived from $\text{Ly}\alpha$ -only line lists and various extended line lists covering additional spectral regions and including other absorption lines. The solid line and solid stars show, respectively, the results from the $\text{Ly}\alpha$ -only MC and real line lists as in Fig. 8. The blue dots and open stars show the corresponding improved σ_v values that result from the addition of the $\text{Ly}\beta$ forest. For the real line lists, the open squares show the effect of further adding the available metal lines in the $\text{Ly}\alpha$ + $\text{Ly}\beta$ forest region. Finally, for two of these QSOs, we show as solid triangles the outcome of using all lines accessible in existing spectra.

est improves σ_v by a factor of 0.88 ± 0.006 . The corresponding comparison for the real line lists (open and solid stars) yields improvement factors of 0.82 to 0.92, with a mean of 0.88, in very good agreement with the MC results.

5.2 Metal lines

In addition to the H I absorption, every high-redshift QSO spectrum shows absorption lines from a number of other ions, such as C IV, Si IV or Mg II. Although these metal lines are far less numerous than the H I lines, they are also much narrower: their widths are of order a few km s^{-1} . In fact, many metal lines are unresolved in current spectra and so the widths of the narrowest lines are unknown. This suggests that metal lines may considerably increase a spectrum's sensitivity to radial velocity shifts and may hence supply valuable additional constraints on \dot{z} .

Could peculiar motions vitiate this supposition? This question is difficult to answer with certainty because the exact origin of many of the various classes of metal absorption lines is still under debate. It is nevertheless evident that the structures responsible for the absorption form an inhomogeneous set (e.g. Churchill et al. 2000, 2007) and represent a large range of environments, depending on the absorbing ion, column density and redshift. For example, low column density C IV absorbers probe the IGM (e.g. Simcoe et al. 2004; Songaila 2006) while strong Mg II lines are associated with galaxies (e.g. Bergeron & Boissé 1991; Steidel, Dickinson & Persson 1994; Zibetti et al. 2005). However, we note that absorbers associated with galaxies are generally found to have absorption cross-sections on the order of tens of kpc (e.g. Steidel 1995; Churchill et al. 2000; Adelberger et al. 2005). For such distances and for halo masses of $\sim 10^{12} M_\odot$ (Steidel et al. 1994; Bouché et al. 2004, 2006) one derives accelerations of a few $\times 10^{-8} \text{ cm s}^{-2}$, while absorber kinematics point to peculiar velocities of no more than a few hundred km s^{-1} (e.g. Churchill et al. 1996). According to Appendix A, peculiar motions

Table 2. Observed lists of metal absorbers.

Name	λ range ^a	N_m^b	Reference
Q1101–264	3050–5765, 5834–8530	225	this work
J2233–606	3400–3850	49	2
HE1122–1648	3500–4091	18	1
HE2217–2818	3550–4050	59	2
HE1347–2457	3760–4335	48	1
Q0302–003	4808–5150	5	1
Q0055–269	4852–5598	14	1
Q0000–26	4300–6450, 7065–8590, 8820–9300	100	this work

References: 1 = Kim et al. (2002), 2 = Kim et al. (2001).

^aWavelength range(s) searched for metal lines in Å.^bNumber of metal absorbers, defined as the number of unique sets of {ion, z , N , b }.

of this magnitude are not problematic. The main caveat here is the unknown fraction of *strong* metal absorbers arising in starburst-driven outflows (e.g. Adelberger et al. 2003; Simcoe et al. 2006; Bouché et al. 2006, see also Section 3.1.2) which must be expected to experience much larger accelerations. However, it may be possible to identify these systems kinematically (Prochaska et al. 2008) and hence to exclude them from the analysis. In any case, in this section we will proceed on the assumption that peculiar motions do not generally invalidate the use of metal lines for a z measurement.

Equation (7) gave a succinct parameterisation of the properties of the Ly α forest. Unfortunately, equivalents for all the various metal line species do not exist, and hence we are unable to gauge the effects of metal lines using MC simulations.

Kim et al. (2001, 2002) and Lu et al. (1996), from whom we obtained the real Ly α forest line lists used in the previous sections, also identified metal lines in their spectra. Unfortunately, they only published measured parameters for lines lying in the Ly α forest region. Although these lines will also have transitions elsewhere (which were probably used in the identification process) the published line lists do not provide a complete view of the metal line population outside of this region.

To improve on this situation we went back to the spectra of two of our QSOs, Q1101–264 and Q0000–26, and derived our own metal line lists covering the entire spectral range available to us. In the case of Q1101–264 the spectrum was the same as that used by Kim et al. (2002), whereas for Q0000–26 we used a UVES/VLT spectrum of similar quality as the HIRES/Keck spectrum used by Lu et al. (1996). In the following we will refer to these two line lists as the ‘complete’ metal line lists.

Table 2 summarises the wavelength ranges that were searched for metal lines and the total number of metal absorbers that were found in each QSO spectrum. The simulated spectra are generated in exactly the same way as in Section 5.1, except that we now add all available metal lines, and extend the spectra of Q1101–264 and Q0000–26 to the red limits given in Table 2.

Fig. 10 shows the effect of the metal lines on σ_v . The green open squares show the result of only using the Ly α +Ly β forest region as in Section 5.1 but adding in all the available metal lines in this region. Comparing this to the H I-only results of the previous section (open red stars) we find that the metal lines improve σ_v by factors of 0.58 to 0.99, with a mean of 0.85. We point out that, strictly speaking, the derived σ_v values are only upper limits for the

six line lists taken from the literature because of their incomplete coverage of the Ly β forest region. Indeed, the best improvement is achieved for one of the complete line lists, Q1101–264 (the lowest redshift QSO), which has a particularly rich metal absorption spectrum (cf. Table 2).

Finally, for our two complete line lists we show as solid green triangles the effect of also adding in the accessible spectral regions redwards of the Ly α forest. The additional metal lines further improve the σ_v values of Q1101–264 and Q0000–26 by factors of 0.55 and 0.90, respectively. With only two values it is obviously impossible to draw a firm conclusion regarding the average improvement offered by the metal lines redwards of the Ly α forest. We therefore choose to be conservative and adopt the larger of the two as a typical value.

Summarising Sections 5.1 and 5.2, the above experiments have shown that the normalisation of equation (15) can be reduced by a factor of $0.88 \times 0.85 \times 0.9 = 0.67$ by considering not just the Ly α forest but all available absorption lines, including metal lines, over the entire accessible optical wavelength range down to a QSO’s Ly γ emission line. Hence we now obtain:

$$\sigma_v = 1.35 \left(\frac{S/N}{2370} \right)^{-1} \left(\frac{N_{\text{QSO}}}{30} \right)^{-\frac{1}{2}} \left(\frac{1+z_{\text{QSO}}}{5} \right)^{-1.7} \text{ cm s}^{-1}. \quad (16)$$

6 MULTIPLE EPOCHS

Fundamentally, a redshift drift experiment consists of simply measuring the velocity shift between two spectra of the same QSO(s) taken at two distinct epochs separated by some time interval Δt_0 . This is the view we took in the previous two sections where we determined the fundamental photon-noise limit of the accuracy with which this shift can be measured, and its scaling behaviour. However, in practice this notion is too simplistic. First of all, it implicitly assumes that the total integration time, t_{int} , required to achieve the necessary S/N would be negligible compared to Δt_0 , so that the two epochs of observation are well-defined. As we will see in the next section this assumption is not valid. Secondly, for a variety of reasons it may be desirable to spread the observations more evenly over the whole period Δt_0 instead of concentrating them in just two more or less well-defined epochs at the interval’s endpoints. Hence the question arises how the accuracy of a redshift drift experiment is affected by distributing the total available t_{int} over multiple observing epochs within the interval Δt_0 .

Let us assume then that observations take place at N_e different epochs, where the j th epoch is separated from the first by Δt_j , so that $\Delta t_1 = 0$ and $\Delta t_{N_e} = \Delta t_0$. We can straightforwardly generalise the framework developed in Section 4.5 by turning equation (10) into a continuous equation for the expected normalised flux in the i th pixel at time Δt :

$$S_i(\Delta t) = S_{0i} + \frac{dS_i}{d\lambda} \lambda_i \frac{\dot{v}_i}{c} \Delta t \equiv S_{0i} + m_i \Delta t. \quad (17)$$

The idea is now to fit this linear relation to the observed fluxes S_{ji} at times Δt_j with errors σ_{ji} , yielding an estimate of the slope m_i and hence of \dot{v}_i for each pixel. (Note that S_{0i} is a nuisance parameter. It represents the ‘true’ flux at the first epoch as opposed to the observed value S_{1i} .) The maximum likelihood estimator for m_i is

$$m_i = \frac{\overline{S_i \Delta t} - \overline{S_i} \overline{\Delta t}}{\overline{\Delta t^2} - \overline{\Delta t}^2}, \quad (18)$$

where the bar denotes the weighted average over all epochs:

$$\bar{x} = \frac{\sum_{j=1}^{N_e} x \sigma_{ji}^{-2}}{\sum_{j=1}^{N_e} \sigma_{ji}^{-2}}. \quad (19)$$

The variance of m_i is given by

$$\sigma_{m_i}^2 = \left[\sum_j \sigma_{ji}^{-2} \left(\overline{\Delta t^2} - \overline{\Delta t}^2 \right) \right]^{-1}. \quad (20)$$

With m_i and its variance in place we can write down the equivalent of equation (12):

$$\sigma_{v_i}^2 = \left(\frac{c}{\lambda_i \frac{dS_i}{d\lambda}} \right)^2 \left[\sigma_{m_i}^2 + \frac{m_i^2}{\left(\frac{dS_i}{d\lambda} \right)^2} \sigma_{S'_i}^2 \right], \quad (21)$$

which in turn allows us to compute \dot{v} averaged over all pixels and its error, σ_v , corresponding to equations (11) and (13). Finally, we re-define $\sigma_v \equiv \sigma_v \Delta t_0$. This new version of σ_v now includes the effect of multiple observing epochs and an arbitrary distribution of the total integration time among them. It is straightforward to show that for $N_e = 2$ one recovers exactly the original σ_v of Section 4.5.

We are now in a position to amend equations (15) and (16) to include the effect of multiple epochs. Recall that these scaling relations were derived for the case of $N_e = 2$ and for equal S/N in the spectra of both epochs. Since the variance of the normalised flux in pixel i scales as the inverse of the integration time we now write σ_{ji} as

$$\sigma_{ji}^2 = \sigma_i^2 \frac{0.5}{f_j}, \quad (22)$$

where f_j is the fraction of the total t_{int} used at the j th epoch ($\sum_j f_j = 1$), and σ_i denotes the flux error (in the i th pixel) that one would obtain if half of the total t_{int} were used. Further defining $\Delta t_j \equiv h_j \Delta t_0$, we can re-write equation (20) as:

$$\begin{aligned} \sigma_{m_i}^2 &= \frac{2\sigma_i^2}{\Delta t_0^2} \left\{ 4 \left[\sum_j h_j^2 f_j - \left(\sum_j h_j f_j \right)^2 \right] \right\}^{-1} \\ &= \sigma_{m_i}^2(N_e = 2, f_1 = f_2 = 0.5) g^2(N_e, h_{1\dots N_e}, f_{1\dots N_e}). \end{aligned} \quad (23)$$

The first term above is just the variance of m_i that one obtains in the case of $N_e = 2$ and equal splitting of t_{int} . The second term is a ‘form factor’ that only depends on the distribution of t_{int} within Δt_0 . Again, it is straightforward to show that $g(N_e = 2, f_1 = f_2 = 0.5) = 1$. Since the form factor is the same for all pixels, and since $\sigma_{m_i}^2$ is the dominant term in equation (21), the sought-after modification of the σ_v scaling relation amounts to simply applying the form factor:

$$\begin{aligned} \sigma_v &= 1.35 \left(\frac{\text{S/N}}{3350} \right)^{-1} \left(\frac{N_{\text{QSO}}}{30} \right)^{-\frac{1}{2}} \left(\frac{1 + z_{\text{QSO}}}{5} \right)^{-1.7} \\ &\times g(N_e, f_{1\dots N_e}) \text{ cm s}^{-1}. \end{aligned} \quad (24)$$

Note that the symbol ‘S/N’ now refers to the *total* S/N per object accumulated over all epochs, in contrast to equations (15) and (16) where it referred to the S/N achieved in each of two epochs.

Note also that we have dropped the dependence of the form factor on $h_{1\dots N_e}$ by considering every night within the period Δt_0 as a potential epoch. This fixes N_e and $h_j = (j-1)/(N_e-1)$, while f_j is constrained to lie in the range $0 \leq f_j \leq l/t_{\text{int}}$, where l is the length of a night (which we will assume to be 9 h on average). Thus we find:

$$g(N_e, f_{1\dots N_e}) = \frac{N_e - 1}{2} \left[\sum_{j=1}^{N_e-1} j^2 f_j - \left(\sum_{j=1}^{N_e-1} j f_j \right)^2 \right]^{-\frac{1}{2}}. \quad (25)$$

Clearly, f_j will be 0 for most nights. Nevertheless, there are obviously a large number of different possible distributions for the f_j s. The best distributions are those that are symmetric and peaked towards the endpoints of Δt_0 . A flat distribution, with equal observations taking place on n equally spaced nights, results in $g = \sqrt{3(n-1)/(n+1)} \approx 1.7$ for $n \gg 1$. Thus the otherwise quite desirable arrangement of observing at a more or less constant rate throughout the period Δt_0 comes with a rather severe penalty attached. A priori, it is difficult to estimate the best g value that can be realistically achieved in practice. From now on we will assume, perhaps somewhat arbitrarily, that all observations occur as much as possible towards the beginning and end of Δt_0 with the constraint that the observing rate averaged over some intermediate timescale of, say, a month cannot exceed $1/3$, i.e. that no more than a third of any month’s telescope time is used for the redshift drift experiment. Depending on the ratio of t_{int} and Δt_0 this results in g values of ~ 1.1 . Essentially, this configuration simply shortens the effective length of the experiment by the amount of time it takes to complete the observations at either end of Δt_0 .

7 CAN WE COLLECT ENOUGH PHOTONS?

In Sections 4, 5 and 6 we learnt what S/N ratio is required to achieve a given sensitivity to radial velocity shifts using QSO absorption spectra. In a photon-noise limited experiment the attainable S/N depends only on four quantities: the brightness of the source, the size of the telescope’s collecting area, the total integration time and the total efficiency. By ‘total efficiency’ we mean the ratio of the number of detected photo-electrons to the number of source photons at the top of the atmosphere, i.e. it comprises atmospheric absorption and all losses occurring in the combined telescope/instrument system, including entrance aperture losses and the detector’s quantum efficiency.

In this Section we will investigate in detail the 5-dimensional parameter space that is spanned by the above four quantities and redshift, in order to determine whether a feasible combination exists that would allow a meaningful \dot{z} measurement.

7.1 S/N formula

We begin by writing down the relation between the S/N per pixel and the above four parameters for the photon-noise limited case:

$$\frac{S}{N} = 700 \left[\frac{Z_X}{Z_r} 10^{0.4(16-m_X)} \left(\frac{D}{42 \text{ m}} \right)^2 \frac{t_{\text{int}}}{10 \text{ h}} \frac{\epsilon}{0.25} \right]^{\frac{1}{2}}, \quad (26)$$

where D , t_{int} and ϵ are the telescope diameter, total integration time and total efficiency, Z_X and m_X are the zeropoint and apparent magnitude of the source in the X -band, respectively, and $Z_r = (8.88 \times 10^{10}) \text{ s}^{-1} \text{ m}^{-2} \mu\text{m}^{-1}$ is the AB zeropoint (Oke 1974) for an effective wavelength of 6170 Å [corresponding to the Sloan Digital Sky Survey (SDSS) r -band]. The normalisation of the above equation assumes a pixel size of 0.0125 Å (see Section 4.4) and a central obscuration of the telescope’s primary collecting area of 10 per cent. $D = 42 \text{ m}$ corresponds to the Baseline Reference Design for the European ELT (E-ELT; Gilmozzi & Spyromilio 2007).

7.2 High-redshift QSOs

The photon flux from QSOs is of course not a free parameter that can be varied at will. Instead we will have to content ourselves

with what will be offered by the population of real QSOs known at the time of a hypothetical z experiment. Here we do not wish to speculate on possible future discoveries of QSOs and hence we will restrict ourselves to the ones known already today. In the following we will extract a list of potential targets for a z experiment from existing QSO catalogues. For each candidate target QSO we will need a reliable magnitude that can be used to estimate its photon flux, as well as its redshift.

The largest QSO catalogue with reliable, homogeneous photometry and redshifts currently available is the fourth edition of the SDSS Quasar Catalog (Schneider et al. 2007). Being based on the fifth data release of the SDSS, it yields 16 913 QSOs with $z_{\text{QSO}} \geq 2$. The catalogue provides PSF magnitudes in the *ugriz* bands which we do *not* correct for Galactic extinction (as is appropriate for S/N calculations). Since we are interested in the continuum flux we will use, for each QSO, the magnitude of the bluest filter that still lies entirely redwards of the QSO's Ly α emission line. Specifically, for objects with $z_{\text{QSO}} < 2.2$ we will use the *g*-band magnitude; for $2.2 \leq z_{\text{QSO}} < 3.47$ the *r*-band; for $3.47 \leq z_{\text{QSO}} < 4.61$ the *i*-band; and for $4.61 \leq z_{\text{QSO}}$ the *z*-band. We then apply a small correction to the selected magnitude to transform the observed flux to that expected at the centre of the Ly α forest assuming a power-law spectral shape of the form $f_\nu \propto \nu^{-0.5}$ (Francis 1993).

Unfortunately, the SDSS catalogue does not cover the whole sky. The largest collection of QSOs covering the *entire* sky is the 12th edition of the Catalogue of Quasars and Active Nuclei recently published by Véron-Cetty & Véron (2006), which contains many additional bright QSOs not included in the SDSS catalogue. However, since the Véron catalogue is a compilation of data from many different sources its photometry is very inhomogeneous and cannot readily be converted to a photon flux. We will dispense with this inconvenience by matching the Véron catalogue to the all-sky SuperCOSMOS Sky Survey (SSS; Hambly et al. 2001). Although the photographic SSS photometry is not endowed with reliable absolute calibrations either, at least it is homogeneous and covers three bands (*B_J*, *R* and *I*), allowing us to synthesise approximate SDSS magnitudes.

We proceed by first selecting all QSOs from the Véron catalogue with $z_{\text{QSO}} \geq 2$ and not flagged as unreliable, resulting in 21 895 objects. For each of these we then identify the nearest SSS object, allowing a maximum separation of 5 arcsec, and retrieve the corresponding SSS catalogue data. 938 Véron objects have no SSS counterpart, presumably because their coordinates are unreliable. We then use those 11 744 objects that are in common to the SDSS and combined Véron-SSS catalogues to derive linear relations between the SDSS and SSS magnitudes, allowing for zeropoint offsets and colour terms. Such relations are reasonable representations of the data and we find that the distributions of the residuals have r.m.s. values of less than 0.3 mag in all cases. Finally, we purge the common objects from the Véron-SSS catalogue and use the above relations to synthesise an SDSS magnitude in the appropriate band (see above) for each of the remaining QSOs.

For those QSOs in the initial Véron catalogue that have no match in the SSS, or which are missing an SSS band needed to synthesise the required SDSS magnitude, we will simply use the *V* or *R*-band magnitude as listed in the Véron catalogue, provided it is non-photographic.

In summary, the final combined sample of 25 974 QSOs is constructed from three sub-sets: (i) SDSS; (ii) objects with redshifts from the Véron catalogue and photometry from the SSS (converted to the SDSS system); and (iii) objects where both the redshifts and




Figure 11. The dots show the known, bright, high-redshift QSO population (separated by sub-sets as indicated, see text) as a function of redshift and estimated photon flux at the centre of the Ly α forest. Along the right-hand vertical axis we have converted the photon flux to a corresponding Johnson *V*-band magnitude. The background colour image and solid contours show the value of σ_v that can be achieved for a given photon flux and redshift, assuming $D = 42$ m, $\epsilon = 0.25$ and $t_{\text{int}} = 2000$ h. The contour levels are at $\sigma_v = 2, 3, 4, 6, 8$ and 10 cm s $^{-1}$. The dotted contours show the same as the solid ones, but for $D = 35$ m or, equivalently, for $\epsilon = 0.17$ or $t_{\text{int}} = 1389$ h.

the photometry are taken from the Véron catalogue. We remind the reader that the quality and reliability of the photometry decreases rapidly from (i) to (iii).

7.3 Achievable radial velocity accuracy

In Fig. 11 we plot our QSO sample, split by the above sub-sets, in the $N_{\text{phot}}-z_{\text{QSO}}$ plane, where N_{phot} is a QSO's photon flux at the top of the atmosphere and at the centre of the QSO's Ly α forest, as implied by the appropriate magnitude described above. Using equations (24) and (26) and assuming values for D , ϵ and t_{int} we can calculate, for any given combination of N_{phot} and z_{QSO} , the value of σ_v that would be achieved if *all* of the time t_{int} were invested into observing a single QSO with the given values of N_{phot} and z_{QSO} . The background colour image and solid contours in Fig. 11 show the result of this calculation, where we have assumed $D = 42$ m, $\epsilon = 0.25$ and $t_{\text{int}} = 2000$ h. Note that we have included both the improvement of σ_v afforded by the Ly β forest and the metal lines as well as the deterioration caused by spreading t_{int} over a 0.9 yr period at either end of a $\Delta t_0 = 20$ yr interval.

From Fig. 11 we can see that, although challenging, a reasonable measurement of $\dot{z}(z)$ is within reach of a 42-m telescope. There exist a number of QSOs that are bright enough and/or lie at a high enough redshift to provide reasonable values of σ_v . We find 18 objects at $\sigma_v < 4$ cm s $^{-1}$ and 5 objects at $\sigma_v < 3$ cm s $^{-1}$, with good coverage of the redshift range 2–4. One object even gives $\sigma_v = 1.8$ cm s $^{-1}$. However, for a smaller telescope with $D = 35$ m the number of objects with $\sigma_v < 4$ cm s $^{-1}$ reduces to only 7 (cf. dotted contours).

In Fig. 12 we show more comprehensively how the number of QSOs with σ_v smaller than a given value depends on the telescope




Figure 12. The colour image and the contours show the number of QSOs for which the σ_v value on the ordinate or better can be achieved for a given combination of telescope size, efficiency and integration time. The contour levels are at $N_{\text{QSO}} = 3, 5, 10$ and 30 .

parameters and integration time, which we summarise into a single ‘normalised observational setup parameter’, O :

$$O = \left(\frac{D}{42 \text{ m}} \right)^2 \frac{\epsilon}{0.25} \frac{t_{\text{int}}}{2000 \text{ h}} \quad (27)$$

(cf. equation 26). For a given value of O the colour image and contours in Fig. 12 show the number of QSOs that would give a σ_v equal to or smaller than the value along the ordinate if all of t_{int} was spent on any one of them. For example, if we wanted to be able to choose our targets from 30 QSOs bright enough and/or at a high enough redshift to be able to achieve $\sigma_v = 3 \text{ cm s}^{-1}$ or better on each object individually, then we would require $O \approx 2.1$. Note however, that the ordinate of Fig. 12 does *not* give the overall value of σ_v for a \dot{z} experiment using the best N_{QSO} targets and setup O . The reason is of course that the total value of σ_v of such an experiment depends on how the total integration time is split up among the N_{QSO} targets (see below).

7.4 Target selection and observing strategies

The question of how to select targets for a \dot{z} experiment depends on what exactly one wishes to achieve. For example, we could simply rank the candidate targets by their achievable σ_v for a fixed observational setup O , as suggested by Fig. 11, and select the objects with the smallest values. This would be the correct selection strategy if the goal were to obtain the smallest possible overall σ_v of the combined sample. Another possible goal might be to maximise the significance with which $\dot{z} \neq 0$ can be detected. In this case we would select the targets with the largest values of $|\dot{v}|/\sigma_{\dot{v}}$, which would give much larger weight to the high redshift objects and would practically deselect all objects at $z_{\text{QSO}} \lesssim 2.5$ where the redshift drift is too small to be detected. Yet another option would be to maximise the sensitivity to Ω_Λ (at a fixed, given value of Ω_M) by selecting the objects with the largest values of $\frac{d\dot{v}}{d\Omega_\Lambda}/\sigma_{\dot{v}}$. It is clear that the choice of selection method will also depend on the possible adoption of priors from other cosmological observations.

Once the exact goal and hence the target selection strategy has been defined we must also choose the number of targets to include

in the experiment. Including more than just the ‘best’ target (according to a given selection strategy) is desirable for several practical reasons, including the ability to observe for as large a fraction of the total available telescope time as possible (to ease scheduling), and to be able to identify any potential directional systematics caused by the atmosphere, the telescope, the instrument, the calibration procedures or the transformation of redshifts to the cosmological reference frame. Obviously though, the more objects are included into the experiment the worse the final result will be (in the absence of systematics) because observing time will have to be redistributed from the ‘best’ object to the less suited ones.

Finally, we must also decide how to divide the total available integration time among the objects included in the experiment. Again, there are several possibilities, including: the time could be split equally among the targets or in such a way as to ensure equal S/N, equal σ_v or an equal value of the selection parameter. (Note that in practice it is likely that operational and scheduling constraints would limit the available choices.)

From the above it is clear that there are a large number of different possibilities of implementing a \dot{z} experiment. In particular, the decision on what precisely the objective of such an experiment should be does not appear to be straightforward. We will now focus on three alternatives which may be perceived to be representative of three different approaches.

1. The simplest possible goal is to aim for the most precise \dot{z} measurement possible, i.e. the smallest overall σ_v . This may be considered a ‘pure experimentalist’ approach where virtually no prior observational information or theoretical expectation is used to (mis-)guide the design of the experiment.

2. Another approach is to emphasise the most basic and yet most unique (and perhaps most captivating) feature of the redshift drift experiment, which is being able to *watch*, literally and in real time, the Universe changing its rate of expansion. In this case the aim is to prove the existence of a dynamical redshift drift effect and hence to measure a non-zero value of \dot{z} with the highest possible significance.³ However, selecting targets according to this approach requires the adoption of, and hence dependence on, a specific model for $\dot{z}(z)$, and is therefore not free of prior assumptions, in contrast to our first example.

3. As explained in the introduction, the discovery and the unknown physical source of the acceleration of the universal expansion provide a strong motivation for any observation seeking to determine $H(z)$, and this is also true for the direct and dynamical method of measuring \dot{z} . From this point of view the goal should be to place the strongest possible constraints on the acceleration. Translating this goal to a target selection strategy again requires a model of $\dot{z}(z)$, including the acceleration.

7.5 Simulations of example measurements and constraints on model parameters

We now proceed to implement each of the above approaches on the QSO sample shown in Fig. 11 in order to illustrate what can be achieved with each of them. However, we stress that the above

³ The value of this approach is perhaps best appreciated by picturing the PI of a future \dot{z} experiment angrily muttering “And yet it moves!” under her/his breath as s/he leaves a seminar room after having tried to convince a sceptical audience that (i) an effect has indeed been observed and (ii) that it is not an instrumental artifact.




Figure 13. The colour image and the contours show the final, overall value of σ_v achieved by targeting the N_{QSO} objects with the best individual σ_v values (cf. Fig. 11) and by employing a given combination of telescope size, efficiency and total integration time. The total integration time is split equally among the targets. The contour levels are at $\sigma_v^{\text{tot}} = 2, 3, 4$ and 5 cm s^{-1} .

approaches and the specific implementations below are only example strategies, and that many variants, refinements and completely different procedures are possible.

We begin with the simplest case of selecting targets by the value of their achievable σ_v , which can be read off Fig. 11 for each object. Having chosen the best N_{QSO} targets by this criterion we can compute the total, overall value of σ_v that can be achieved for this sample and for a given observational setup, O :

$$\sigma_v^{\text{tot}} = \left[\sum_i^{N_{\text{QSO}}} \sigma_v^{-2}(z_{\text{QSO},i}, m_i, f_{\text{obj},i} O) \right]^{-\frac{1}{2}}, \quad (28)$$

where $f_{\text{obj},i}$ is the fraction of the total integration time allocated to the i th object. In Fig. 13 we plot σ_v^{tot} as a function of N_{QSO} and O , where we have chosen to distribute the time equally among all objects (i.e. $f_{\text{obj},i} = N_{\text{QSO}}^{-1}$). Distributing the time to give equal S/N instead (which may be operationally desirable) leads to slightly better values of σ_v^{tot} when only a few QSOs are targeted. The reason is of course that the object with the best σ_v is only the seventh brightest target (cf. Fig. 11) and so it is allocated relatively more time. However, for $N_{\text{QSO}} \gtrsim 15$ this advantage is lost and equal time allocation produces σ_v^{tot} values that are smaller by ~ 10 per cent compared to equal S/N allocation.

From Fig. 13 we can see that an overall value of $\sigma_v^{\text{tot}} \approx 2\text{--}3 \text{ cm s}^{-1}$ is well within reach of an ELT, even when tens of objects are targeted for the experiment. The steepness of the contours also indicate that σ_v^{tot} does not depend very sensitively on N_{QSO} , at least for $N_{\text{QSO}} \gtrsim 10$.

For further illustration we now perform a MC simulation of a redshift drift experiment using this target selection strategy. We (arbitrarily) choose $N_{\text{QSO}} = 20$ and $O = 2$ so that the overall accuracy of the experiment is $\sigma_v^{\text{tot}} = 2.34 \text{ cm s}^{-1}$ (cf. Fig. 13). Furthermore we assume that the observations span a time interval of $\Delta t_0 = 20 \text{ yr}$ and that the true $\dot{z}(z)$ is given by our standard cosmological model. The result is presented in Fig. 14 where we show as blue dots one realisation of the ‘observed’ velocity drifts




Figure 14. The three sets of ‘data’ points show MC simulations of the redshift drift experiment using the three different example implementations discussed in the text. In each case we have assumed an observational setup of $O = 2$ and we plot as ‘data’ points the ‘observed’ values and error bars of the velocity drift \dot{v} , expected for a total experiment duration of $\Delta t_0 = 20 \text{ yr}$ and for standard cosmological parameters ($h_{70} = 1$, $\Omega_M = 0.3$, $\Omega_\Lambda = 0.7$). For a given QSO we use the centre of the Ly α forest as the redshift of the \dot{v} measurement. Blue dots: selection by σ_v , $N_{\text{QSO}} = 20$ (binned into four redshift bins), equal time allocation. Yellow squares: selection by $|\dot{v}|/\sigma_{\dot{v}}$, $N_{\text{QSO}} = 10$ (in two redshift bins), equal time allocation. Brown triangles: selection by best combined constraint on Ω_Λ , $N_{\text{QSO}} = 2$, optimal time distribution. The solid lines show the expected redshift drift for different parameters as indicated, and $h_{70} = 1$. The grey shaded areas result from varying H_0 by $\pm 8 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

and their errors along with the input model (red solid line). Since the selected 20 targets cover the redshift range $2.04 \leq z_{\text{QSO}} \leq 3.91$ quite homogeneously we have binned the measurements into four equally sized redshift bins.

By construction these points represent the most precise measurement of \dot{z} that is possible with a set of 20 QSOs (using equal time allocation) and $O = 2$. However, since many of the selected QSOs lie near the redshift where $\dot{z} = 0$ (for the assumed model) the redshift drift is only detected with an overall significance of $S = |\bar{\dot{v}}|/\sigma_{\dot{v}}^{\text{tot}} = 1.4$, where $|\bar{\dot{v}}|$ is the weighted mean of the absolute values of the expected velocity drifts.

This can be improved upon by choosing our second approach and selecting targets by the largest value of $|\dot{v}|/\sigma_{\dot{v}}$. This quantity is a strongly increasing function of redshift and this selection strategy results in quite a different set of objects: the best 20 targets according to this criterion include only 3 of the objects previously selected. For $O = 2$, $\Delta t_0 = 20 \text{ yr}$ and equal time allocation as before, the combined overall significance of this sample is $S = 2.7$ (assuming our standard model of \dot{z} is correct). The yellow squares in Fig. 14 show the result of a MC simulation using this setup and implementation, except that we use $N_{\text{QSO}} = 10$ (in two redshift bins) which gives a slightly better significance of $S = 3.1$. To reach $S \geq 4$ we need to further reduce N_{QSO} to 3, or instead increase O to 3.4 or Δt_0 to 25.4 yr.

Finally, we turn to our third approach and the question of how to best select targets to constrain the acceleration of the expansion and what can be achieved in this respect with our sample of known QSOs. As in the previous case the answer will depend on what to expect for the expansion history and in particular for the acceleration. For the purpose of the following discussion we will again assume our standard cosmological model.

The simplest thing we can do to constrain the acceleration is to unambiguously detect its existence, i.e. to measure $\dot{z} > 0$ with the highest possible significance. This implies (i) that we need a \dot{v} measurement at $z < z_0$, where z_0 is defined by $\dot{v}(z_0) = 0$, and (ii) that target selection should proceed by the largest value of \dot{v}/σ_v , which indeed favours the lowest available redshifts. However, even if we use only the single best object by this criterion and assume a generous $\Delta t_0 = 25$ yr then a $2\sigma_v$ detection of $\dot{v} > 0$ would still require an unfeasible $O = 9.5$. The reason for not being able to do better is of course our inability to access the Ly α forest at $z \approx 0.7$ where \dot{v} is the largest.

Let us analyse the situation more systematically by switching to the parameter space of our cosmological model (which is 3-dimensional since we will *not* assume spatial flatness). In this parameter space our goal of detecting the acceleration translates to proving that the deceleration parameter $q_0 = \frac{\Omega_M}{2} - \Omega_\Lambda$ is < 0 . However, in our model the acceleration is due to a cosmological constant, and so for simplicity we will instead pursue the slightly easier goal of proving the existence of a cosmological constant, i.e. of placing a positive lower limit on Ω_Λ after having marginalised over Ω_M and H_0 .

Consider the constraint of a single measurement of $\dot{v} = \Delta v / \Delta t_0$ at some z (i.e. provided by some QSO), which happens to be precisely equal to its expected value. Obviously, a single data point cannot constrain all three parameters at the same time but only some 2-dimensional surface. Solving equations (4) and (5) for Ω_Λ we find that this surface is defined by:

$$\Omega_\Lambda = \frac{(1+z)^2}{(1+z)^2 - 1} \left[z\Omega_M + 1 - \left(1 - \frac{\dot{v}}{cH_0} \right)^2 \right]. \quad (29)$$

Since \dot{v} is assumed to be equal to its expectation value this constraint surface is entirely defined by z and it must include the point $(\Omega_M, \Omega_\Lambda, h_{70}) = (0.3, 0.7, 1)$. Furthermore, the measurement error σ_v endows the surface with some small thickness. Hence we can see that each QSO in our sample [representing a (z, σ_v) pair] defines its own constraint surface. The question is whether there exists some combination of some sub-set of these surfaces such that the projection of the combined constraint onto the Ω_Λ -axis does not include the region $\Omega_\Lambda \leq 0$. If so, what is the largest significance level at which $\Omega_\Lambda = 0$ can be excluded?

To address this problem let us fix $h_{70} = 1$ for now, so that a given constraint surface turns into a band in the Ω_Λ - Ω_M plane (cf. the grey shaded bands in Fig. 15). Consider first the case $N_{\text{QSO}} = 1$. For a fixed H_0 , equation (29) simply gives a linear relation between Ω_Λ and Ω_M which has a non-zero slope for all $z > 0$. It follows that the projection of the constraint band onto the Ω_Λ -axis will be infinite unless we can place some kind of limit on Ω_M . Since the slope is positive we require a lower limit on Ω_M in order to obtain a lower limit on Ω_Λ . Rather uncontroversially we can impose $\Omega_M > 0$. Applying this limit to equation (29) we find that a single \dot{v} measurement can place a positive lower limit on Ω_Λ if the intercept of the Ω_Λ - Ω_M relation, b , is significantly larger than 0. The error on b is given by the width of the constraint band along the Ω_Λ -direction:

$$\Delta b = \Delta\Omega_\Lambda = \frac{d\Omega_\Lambda}{d\dot{v}} \sigma_v = \frac{2}{cH_0^2} \frac{(1+z)H(z)}{(1+z)^2 - 1} \frac{\sigma_v}{\Delta t_0}. \quad (30)$$

and the object giving the best lower limit on Ω_Λ is the one with the largest value of $b/\Delta b$. However, we find that even the best object by this criterion does not deliver a significant lower limit on Ω_Λ . Hence we find that not only are we unable to detect $\dot{v} > 0$ with any




Figure 15. Expected constraints in the Ω_Λ - Ω_M plane from two measurements of the redshift drift at two different redshifts as indicated, assuming $O = 2$ and $\Delta t_0 = 20$ yr. The two targets and the split of the observing time among them were chosen to give the best lower limit on Ω_Λ . The red and blue solid lines and the grey shaded bands show the individual constraints provided by each of the two objects assuming a fixed $h_{70} = 1$. The coloured ellipses show the joint 68 and 90 per cent confidence regions that result from combining the two measurements, marginalising over H_0 using an external prior of $H_0 = (70 \pm 8) \text{ km s}^{-1} \text{ Mpc}^{-1}$. The hashed region indicates the 95 per cent lower limit on Ω_Λ . The dashed contour shows the 68 per cent confidence region that results from using a flat prior on H_0 and just marginalising over the range $0 \leq H_0 \leq 140 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Flat cosmologies and the boundary between current de- and acceleration are marked by solid black lines. The dark shaded region in the upper left corner designates the regime of ‘bouncing universe’ cosmologies which have no big bang in the past.

significance but it is also impossible to prove $\Omega_\Lambda > 0$ with a *single* \dot{v} measurement.

However, this conclusion rests entirely on the weak lower limit on Ω_M that we used ($\Omega_M > 0$). The larger this limit the stronger also the lower limit on Ω_Λ . In other words, a given \dot{v} measurement translates to a different constraint on Ω_Λ depending on what we assume for Ω_M . Hence, even if we cannot detect $\dot{v} > 0$ with high significance it is nevertheless possible to place a significant lower limit on Ω_Λ . All we need is a stronger constraint on Ω_M .

The point is that this constraint can be supplied by the \dot{z} experiment itself, without having to resort to an external prior, by means of a second measurement at *high* redshift (as opposed to at $z \lesssim z_0$). The idea is to combine the measurements from two objects (or two groups of objects) that are selected in different ways, such that their constraint bands in the Ω_Λ - Ω_M plane make as large an angle as possible while also being as narrow as possible (cf. Fig. 15). The first group is selected to have shallow constraint lines, implying low redshifts ($z \lesssim z_0$) and supplying a lower (but by itself insufficient) limit on Ω_Λ as discussed above, while the second group is selected to have steep constraint lines, implying high redshifts (cf. equation 29) and supplying a lower limit on Ω_M . This is illustrated in Fig. 15 where the red and blue lines and the grey shaded areas show the constraint bands from two \dot{v} measurements at low and high redshifts as indicated. If the angle between them is large enough and if the individual σ_v s are small enough then the combination of the

two constraints (in a χ^2 sense) results in an ellipse in the Ω_Λ - Ω_M plane whose lower edge excludes $\Omega_\Lambda = 0$.

Hence we find that we need at least $N_{\text{QSO}} = 2$. However, recall that so far we have kept H_0 fixed. Let us now reconsider the above as a function of H_0 . First of all, we note that the constraint line resulting from the \dot{z} measurement at $z \approx z_0$ is very insensitive to the value of H_0 (because $\dot{z} \approx 0$). Secondly, for the high redshift measurement the constraint line shifts upwards in Fig. 15 as H_0 increases. Hence the centre of the joint ellipse (where the constraint lines cross) moves from $(\Omega_M, \Omega_\Lambda) = (0.3, 0.7)$ approximately along the low- z constraint line towards the lower left of Fig. 15. However, at the same time the extent of the ellipse decreases because a larger H_0 predicts a larger redshift drift and hence results in stronger constraints on Ω_M and Ω_Λ for fixed measurement errors. The net result is that the lower edge of the joint ellipse and hence the obtainable lower limit on Ω_Λ are relatively insensitive to H_0 . As we will see more quantitatively below, a third \dot{z} measurement to constrain H_0 is therefore not necessary.

Thus we find that the strategy of selecting two groups of targets as described above should provide the best possible lower limit on Ω_Λ , even after marginalising over H_0 . However, when trying to express this approach in equations in order to proceed with the target selection one realises that the selection of the two groups cannot in fact be separated, making an analytical procedure impractical. Hence we resort to the simplest case of only one object per group and simply try all possible combinations of two objects in our sample. For any given combination we also determine empirically the optimal split of the total integration time among the two objects. In this way we find that the best lower limit on Ω_Λ is achieved for the two very bright objects in Fig. 11 at $z_{\text{QSO}} = 2.15$ and 3.91 , with the latter object receiving 0.3 of the total integration time. Mock \dot{z} measurements for these objects are shown as brown triangles in Fig. 14.

In Fig. 15 we present the expected constraints from these two objects in the Ω_Λ - Ω_M plane, assuming $O = 2$ and $\Delta t_0 = 20$ yr. The red and blue solid lines show the individual constraint lines for a fixed $h_{70} = 1$ (equation 29) and the grey shaded areas around them are the error bands that correspond to the respective values of σ_v for each object (equation 30). Note that the optimal time split (in the sense of giving the best lower limit on Ω_Λ) is such that the two bands are of roughly equal width.

The coloured ellipses in Fig. 15 show the joint 68 and 90 per cent confidence regions that result from the combination of the two \dot{z} measurements. Here we have assumed that an external constraint places 1σ limits of $\pm 8 \text{ km s}^{-1} \text{ Mpc}^{-1}$ on H_0 (Freedman et al. 2001) before marginalising over this parameter. By comparing the error bars on the brown triangles in Fig. 14 with the grey shaded region around the red line we can see that this uncertainty in H_0 is small relative to the errors on \dot{z} . Hence the ellipses in Fig. 15 are almost those that would be obtained for a fixed H_0 .

Projecting the joint constraints onto the vertical axis we find that the 95 per cent lower limit on the cosmological constant is $\Omega_\Lambda > 0.16$ (indicated by the hashed region in Fig. 15), while $\Omega_\Lambda = 0$ is excluded at the 98.2 per cent confidence level. Note that we can also exclude $q_0 = 0$ at a similar level of 97.6 per cent. These numbers do not depend sensitively on the adopted prior on H_0 , for reasons explained above. Doubling the error on H_0 to $\pm 16 \text{ km s}^{-1} \text{ Mpc}^{-1}$ has almost no effect at all. Even if we use a flat prior and simply marginalise over the range $0 \leq H_0 \leq 140 \text{ km s}^{-1} \text{ Mpc}^{-1}$ we can still reject $\Omega_\Lambda = 0$ at 98.1 per cent confidence, although the 95 per cent lower limit now reduces to 0.08 . For a shorter experiment duration of $\Delta t_0 = 15$ yr (and the




Figure 16. Expected constraints in the Ω_Λ - H_0 plane for flat cosmologies assuming $O = 2$ and $\Delta t_0 = 20$ yr and using the same two redshift drift measurements as in Fig. 15. The red and blue solid lines and the grey shaded bands show the individual constraints provided by each of the two objects. The coloured ellipses show the joint 68 and 90 per cent confidence regions that result from combining the two measurements.

flat H_0 prior) we can still exclude $\Omega_\Lambda = 0$ at 93.5 per cent confidence.

Thus we conclude that, assuming general relativity to be valid, an ELT will enable us to confirm the existence of a cosmological constant using purely dynamical evidence – by measuring the redshift drift in QSO absorption lines over a period of ~ 20 yr – *without assuming flatness, using any other external cosmological constraints or making any other astrophysical assumptions whatsoever.*

Combining the constraints of a \dot{z} experiment with those from other cosmological observations would of course lead to significantly stronger results. A detailed examination of this issue is, however, beyond the scope of this paper. We will let it suffice to make two general points: (i) the \dot{z} constraints will be unique and therefore new in the sense that they cannot be provided by any other cosmological observation. \dot{z} probes $H(z)$ in a redshift range that is inaccessible to other methods (such as SNIa, weak lensing, BAO) and this will remain true even in the ELT era. In terms of characterising $H(z)$, \dot{z} measurements at $z \gtrsim 2$ will nicely complement the data at $z \lesssim 2$ provided by SNIa surveys. In the Ω_Λ - Ω_M plane these datasets offer similar but nevertheless distinct constraints. (ii) The datasets that are most complementary to the redshift drift in the Ω_Λ - Ω_M plane are those that constrain the geometry of the Universe, such as the fluctuation power spectrum of the CMB, which can break the remaining degeneracy in the \dot{z} constraints. In fact, the WMAP 3-year data constrain the quantity $\Omega_\Lambda + 0.72 \Omega_M$ (Spergel et al. 2007), so that the CMB degeneracy line is almost exactly orthogonal to the \dot{z} constraints shown in Fig. 15. Combining these constraints would lead to individual 2σ errors in Ω_Λ and Ω_M of ~ 0.08 and ~ 0.06 , respectively.

Finally, let us consider models in which the Universe is *exactly* flat ($\Omega_k = 0$). The parameter space of these models is 2-dimensional and a single \dot{z} measurement results in a constraint line given by

$$\Omega_\Lambda = \frac{(1+z)^2}{(1+z)^3 - 1} \left[1 + z - \left(1 - \frac{\dot{v}}{cH_0} \right)^2 \right]. \quad (31)$$

In Fig. 16 we show the constraints in the Ω_Λ - H_0 plane offered by the two \dot{v} measurements considered above ($O = 2$, $\Delta t_0 = 20$ yr). Clearly, H_0 is not constrained in any useful way and both substantially higher and lower redshift measurements of similar or better quality would be required to do so. On the other hand, Ω_Λ is constrained quite well. Marginalising over H_0 we find a 2σ range of $[0.42, 0.74]$, and $\Omega_\Lambda = 0$ is excluded at 98.7 per cent confidence. For a shorter experiment duration of $\Delta t_0 = 15$ yr we can still exclude $\Omega_\Lambda = 0$ at 94.7 per cent confidence. Note that these latter values are not very much higher than in the general case above. The difference is of course that in the flat case the best fit along the $\Omega_\Lambda = 0$ axis is obtained for an unfeasibly low H_0 value of $h_{70} = 0.1$, while in the general case this value is 1. Hence, for flat cosmologies any external H_0 prior has a strong impact on the lower limit one can place on Ω_Λ , in contrast to the general case. For example, applying the relatively weak prior that H_0 is known to within $\pm 16 \text{ km s}^{-1} \text{ Mpc}^{-1}$ results in the rejection of $\Omega_\Lambda = 0$ at > 99.99 per cent confidence for any $\Delta t_0 > 9$ yr or, alternatively, for $\Delta t_0 = 15$ yr and any $O > 0.5$.

7.6 Constraints on alternative cosmological models

So far we have only considered the standard cosmological model, i.e. general relativity with a stress-energy tensor dominated by cold (dark) matter and a cosmological constant (Λ CDM). Needless to say, industrious theorists have proposed numerous alternative explanations for the observed acceleration of the universal expansion. Broadly speaking, these models either suggest alternative forms of dark energy or modify the theory of gravity. In Appendix B we give a brief overview of two examples for each of these two classes of models. These are quintessence, generalised Chaplygin gas, Cardassian expansion and the DGP brane-world model. Despite fundamental differences, some generic features are shared by all four. Obviously, they all ‘predict’ late-time acceleration. They also agree that the early evolution of the Universe is indistinguishable from matter-dominated expansion in general relativity (so as not to spoil the successes of the standard model in explaining early structure formation). Furthermore, all of them have at least as many free parameters as Λ CDM. In fact, three of the four models considered here include Λ CDM phenomenologically as a special case, i.e. as a sub-space of the models’ parameter space. With these properties in mind, and given that Λ CDM can successfully reproduce SNIa, CMB and BAO data, it is perhaps not too surprising that the alternative models can also be adjusted to fit these datasets, resulting in best-fit parameter values consistent with Λ CDM (e.g. Davis et al. 2007).

Given this state of affairs we do not think it justified to study these alternative models, in this paper, to the same level of detail as we did for Λ CDM in the previous section. The name of the game would be to ask, for each model, how much of the non- Λ CDM-like parameter space could be excluded by a redshift drift experiment. This would best be done by considering the \dot{z} constraints jointly with those from other relevant datasets. Instead, we will do something much simpler here. As already mentioned above, Davis et al. (2007) have fit all four of our alternative models to current SNIa, CMB and BAO data. The idea is now to ask, for each model, which values of \dot{z} are still ‘allowed’ given the constraints already placed on the model’s parameters by the Davis et al. analysis. Comparing

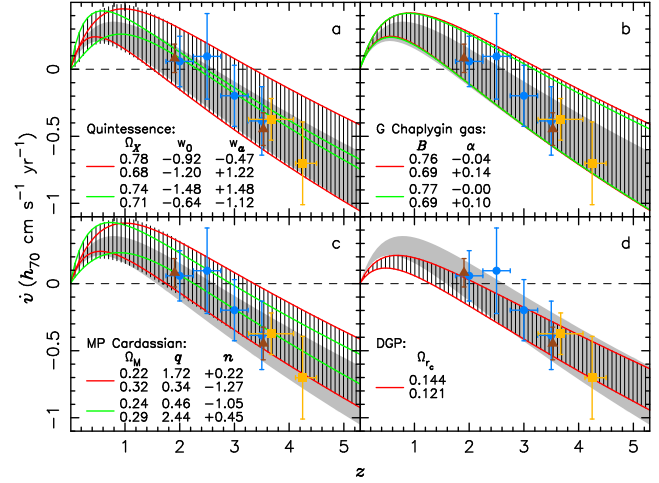


Figure 17. Comparison of the predictions for the redshift drift by non-standard cosmological models. For simplicity, we impose flatness and fix $h_{70} = 1$ in all cases. (a) At a given redshift the vertically hashed region shows the range of \dot{z} values predicted by quintessence, allowing the model’s three free parameters to vary within the joint 95.4 per cent confidence region imposed by SNIa, CMB and BAO data (Davis et al. 2007). The green and red solid lines show the predictions for those parameter combinations that produce the most extreme allowed \dot{z} values at $z = 0.3$ and 4 , respectively. $(w_0, w_a) = (-1, 0)$ corresponds to flat Λ CDM. (b–d) Same as (a) for the other models as indicated. For the Chaplygin gas and Cardassian models $\alpha = 0$ and $(q, n) = (1, 0)$ correspond to flat Λ CDM, respectively. In each panel we also show, for comparison, the predictions of the flat Λ CDM model with $0.65 \leq \Omega_\Lambda \leq 0.74$ as the grey shaded region. We also reproduce the three sets of simulated measurements from Fig. 14.

these allowed ranges with the simulated \dot{z} measurements of the previous section will give a good impression of what can be achieved.

In Fig. 17 we plot this allowed \dot{z} range as the vertically hashed region for all four models considered here along with the predictions for some specific parameter combinations as indicated. Here we have assumed $\Delta t_0 = 20$ yr. In order to reduce the number of free parameters we are only considering flat models with H_0 fixed at $h_{70} = 1$. For each model the plotted \dot{z} range corresponds to the joint 95.4 per cent confidence region in the space spanned by the model’s remaining free parameters. The grey shaded region in each panel shows the flat Λ CDM model (which has only one free parameter since we have fixed Ω_k and H_0), varying Ω_Λ in the range $[0.65, 0.74]$. This corresponds to the 95.4 confidence range obtained from the two redshift drift measurements discussed in the previous section. Note that the hashed and grey shaded regions are not ‘centred’ on each other. The reason is that we chose to centre the grey shaded region on our fiducial value of $\Omega_\Lambda = 0.7$, whereas the ranges of the equivalent parameters in panels (a)–(c) are centred on 0.73. For reference, we also reproduce, in each panel, the simulated data points from Fig. 14.

In panels (a)–(c) the hashed regions are somewhat larger than the grey shaded regions and the error bars on the ‘data’ points, indicating that a redshift drift experiment will be able to significantly reduce the sizes of the allowed regions in these models’ parameter spaces. The specific parameter combinations that we chose for the solid lines are those that produce the most extreme allowed \dot{z} values at $z = 0.3$ (green lines) and at $z = 4$ (red lines). For the quintessence and Cardassian models we can see that these two sets of lines have very different parameter values. Indeed, the two parameter combinations that produce the extreme \dot{z} values at low

redshift predict very similar values at high redshift, and vice versa (although to a lesser extent). This again highlights the complementarity between $H(z)$ constraints at low and high redshifts. Notice, however, that the situation is different in panel (b). In the case of generalised Chaplygin gas one set of parameters produces the maximum (or minimum) allowed \dot{z} values at more or less all redshifts. The reason is that the region in this model's 2-dimensional parameter space allowed by the Davis et al. analysis is in fact nearly 1-dimensional. This is of course exactly true for the DGP model (panel d) which has only one free parameter to begin with. Notice, however, that in this case $\dot{v}(z)$ has a somewhat different slope. In contrast to the other models, the DGP prediction for \dot{z} cannot be made arbitrarily close to the Λ CDM prediction and from Fig. 17 we can see that the differences are of roughly the same size as the expected \dot{z} measurement errors.

We defer a detailed investigation of these issues to a future analysis. Here we simply conclude that the increased redshift range of $H(z)$ measurements afforded by a redshift drift experiment will significantly add to our capability of discriminating between alternative cosmological models.

8 CONCLUSIONS

In this paper we have appraised the prospects of using an ELT to detect and characterise the cosmological redshift drift, \dot{z} , by means of a \sim two decade-long spectroscopic monitoring campaign of high redshift QSO absorption lines. This experiment would be unique among cosmological observations in that it directly probes the global dynamics of the Robertson-Walker metric, thereby offering a novel and truly independent path to the expansion history of the Universe. We summarise our results as follows:

1. We confirm the earlier proposition (Loeb 1998) that the Ly α forest is the most promising target for a redshift drift experiment. We have explicitly examined the peculiar motions of the gas responsible for the Ly α forest in a hydrodynamical simulation and found their effects negligible. We have also assessed the effects of variations of the physical properties of the absorbing gas over the timescale of the experiment and found them similarly insignificant.

2. Under the assumption of photon-noise limited observations the accuracy, σ_v , to which a putative radial velocity shift between two spectra of the same object can be determined depends chiefly on the number and sharpness of the relevant spectral features and the spectra's S/N. The number density of the absorption lines that constitute the Ly α forest depends strongly on redshift. Using extensive Monte Carlo simulations we have hence derived a quantitative relation between the σ_v of a pair of Ly α forest spectra on the one hand, and the spectra's S/N and the background QSO's redshift on the other hand (equation 15).

3. Apart from the Ly α forest, all QSO spectra also display higher order H I lines and a variety of metal absorption lines. Using the H I Ly β lines in addition to the Ly α lines, as well as all available metal lines within and to the red of the Ly α forest significantly improves a given spectrum's sensitivity to radial velocity shifts by a factor of 0.67 (equation 16).

4. In practice it will not be possible to measure the redshift drift by simply comparing just two observations of some set of objects obtained at two well-defined epochs separated by some Δt_0 . In reality, the total observing time required for the experiment will be spread over many nights (= epochs) whose distribution within Δt_0 (now defined as the interval between the first and last observations) will be subject to many practical constraints. In Section 6

we have calculated the relevant 'form factor' by which σ_v increases (or, equivalently, by which the effective Δt_0 decreases) as a result of a given observing time distribution. For realistic distributions we find factors of ~ 1.1 to 1.7 .

6. The detailed implementation of a redshift drift experiment will depend on the exact objectives one hopes to achieve. We have elaborated three different possible top-level goals and we have worked out the specific target selection strategies required to tailor a \dot{z} experiment to each of these. Applying these selection schemes to a comprehensive list of high redshift QSOs *already known today* we use our final σ_v scaling relation to predict an ELT's capability to measure \dot{z} . For example, selecting those 20 targets that deliver the greatest measurement accuracy we find that a 42-m ELT could achieve a total, overall accuracy of 2.34 cm s^{-1} in 4000 h of observing time (Fig. 13). If we additionally assume that our standard cosmological model correctly predicts the size of the redshift drift then we will be able to prove its existence (i.e. measure a non-zero value) at a significance of 3.1σ in $\Delta t_0 = 20 \text{ yr}$ (cf. Fig. 14).

Unfortunately however, due to the fact that the Ly α forest is only accessible from the ground for $z \gtrsim 1.7$ it will *not* be possible to directly measure $\dot{z} > 0$ (the hallmark of past accelerated expansion) from any *one* of the QSOs in our sample with any significance, unless the transition from decelerated to accelerated expansion occurred significantly earlier than predicted by our standard model. However, even if the standard model's prediction for \dot{z} is correct an ELT will nevertheless be able to provide unequivocal proof of the existence of past acceleration by obtaining an upper limit on \dot{z} at high redshift and a lower limit at $z \approx 2$ where the transition from $\dot{z} < 0$ to > 0 is expected to occur (cf. Fig. 14). For the above experiment parameters these limits will be strong enough to prove that \dot{z} must become positive below $z \approx 2$. We stress that the validity of this proof only requires that gravity can be described by a metric theory and that the Universe is homogeneous and isotropic on large scales. It is entirely independent of any specific theory of gravity (such as general relativity) and does not require any assumptions regarding the spatial geometry of the Universe. It is also independent of other cosmological observations and does not rely on any other astrophysical assumptions whatsoever. Hence, a \dot{z} experiment carried out with an ELT will arguably provide the most direct evidence of acceleration possible.

In the context of general relativity dominated by cold dark matter and a cosmological constant [$(\Omega_M, \Omega_\Lambda) = (0.3, 0.7)$] a \dot{z} experiment over $\Delta t_0 = 20 \text{ yr}$ using 4000 h of observing time on a 42-m ELT will be able to exclude $\Omega_\Lambda = 0$ with 98.1 per cent confidence, providing a 2σ lower limit of 0.08 (cf. Fig. 15). Additionally assuming spatial flatness results in a 2σ range of $0.42 < \Omega_\Lambda < 0.74$ (Fig. 16).

7. Finally, we point out that a \dot{z} experiment constrains $H(z)$ over a redshift range that is inaccessible to other methods. Redshift drift measurements at $z \gtrsim 2$ would therefore complement current and future SNIa and BAO surveys. Adding \dot{z} to the suite of modern cosmological tests would thus provide additional leverage in constraining the parameters of both Λ CDM and alternative cosmological models (cf. Fig. 17).

We stress that all the results summarised above were derived under the assumption that the uncertainty on radial velocity shift measurements will be dominated by photon noise. The next step is to understand to what extent this assumption can be realised in practice. To this end we need to (i) draw up a list of all potential sources of error, both astrophysical and instrumental in nature, (ii) quantitatively assess their impact on a \dot{z} measurement, and (iii) where necessary and possible identify appropriate hardware

solutions, calibration and/or analysis techniques that will reduce their impact below the level of the photon noise. Although we have already addressed some of the most obvious astrophysical sources of noise in Section 3.1, a realistic and comprehensive assessment of the total error budget is a non-trivial task – in particular because of point (iii) above – which we defer to a future paper.

One of the potential error sources that will require careful attention is the fact that QSO continua are to some extent variable. The flux-difference method we described in Section 4.5 to extract the redshift drift signal from the data is clearly very sensitive to any errors in determining the change of the continuum. Similarly, it is also quite susceptible to any residual uncertainties in the sky subtraction, flat-fielding and stray light correction. Clearly, a flux-difference type of approach would not be the method of choice for a real-life \dot{z} experiment – a conclusion we had already reached in Section 4.5 for an entirely different reason. However, other methods of extracting the redshift drift signal are likely to be less sensitive to these issues. As already suggested in Section 5 we believe that the most promising approach is to simultaneously model the absorption in all spectra of a given object with Voigt profiles, including the \dot{z} of each component (or a model thereof) as an additional free parameter. The drawback of this method is that a complex absorption feature can usually be modelled in more than one way, none of which necessarily has to correspond to reality, and this could represent an additional source of error for a \dot{z} measurement. In any case, the question of how to deal with QSO continuum variability remains to be investigated.

Another issue we have not yet studied is the effect of the evolution of structure along the line of sight to a distant QSO. Consider two photons whose redshift difference we wish to measure. Along their way to us both will encounter various gravitational potential wells. Any evolution of these potentials between the passage of the first and second photon will cause them to follow slightly different paths, perturbing their redshift difference. This effect will have to be studied with high time resolution simulations of cosmic structure formation.

The quest for photon-noise limited radial velocity shift measurements also places rather exacting requirements on the spectrograph to be used for a redshift drift experiment, in particular with respect to its stability and wavelength calibration. However, an instrument concept study investigating these and other issues has already been carried out (CODEX; Pasquini et al. 2005), concluding that the technical challenges are demanding but not insurmountable. One of the novel concepts identified by this study is the use of a ‘laser frequency comb’ for wavelength calibration. This system provides a series of uniformly spaced, narrow lines whose *absolute* positions are known a priori with relative precision better than 10^{-12} (see Murphy et al. 2007 for details). Note that an absolute wavelength calibration of this precision would allow the combination of data from different ELTs and would guarantee the legacy value of the data by enabling future astronomers to measure \dot{z} over much longer timescales.

In deriving our results we have also used the magnitudes and redshifts of the high-redshift QSO population as it is known today. Note that we have used *all* known QSOs, irrespective of their declination. In exchange, so to speak, we have refrained from speculating on future discoveries of bright, high-redshift QSOs. It is clear that existing QSO catalogues do not yet cover the entire sky uniformly with the same completeness, even at the bright magnitudes we are interested in here. Indeed, only one third of the objects shown in Fig. 11 lie in the south. Furthermore, since colour selection is one of the most popular methods of QSO candidate

identification it is likely that our current sample is incomplete in the redshift range $2.4 \lesssim z \lesssim 3$ where the colours of QSOs are very similar to those of main sequence stars (e.g. Richards et al. 2002). Hence it is not unreasonable to believe that at least a few bright QSOs still remain to be discovered, especially in the south. Considering that a \dot{z} experiment would benefit greatly from *any* new discoveries we propose that a complete, systematic survey for bright, high-redshift QSOs be carried out once the site for an ELT has been chosen.

In addition, we propose that all potential target QSOs should be monitored photometrically. Not only would this remove the uncertainties of our present photometry (note that we had to rely on photographic photometry from the SuperCOSMOS Sky Survey for most of the brightest objects in Fig. 11) but it would also allow us to account for QSO variability in the target selection. If the monitoring frequency were high enough (as can be expected from possible future facilities like the Large Synoptic Survey Telescope) then these data could also be used to time the spectroscopic observations of a given target to coincide with a peak in the object’s lightcurve.

We close by observing that any QSO spectra collected as part of a \dot{z} experiment could also be used to address numerous other important issues, including the variability of the fine-structure constant and the primordial abundance of deuterium. Given the substantial observing time requirements for a \dot{z} experiment it is clear that these synergies will play an important role in any realistic attempt to establish a future ELT redshift drift campaign.

ACKNOWLEDGEMENTS

Numerical computations were performed on the COSMOS supercomputer at the Department of Applied Mathematics and Theoretical Physics in Cambridge. COSMOS is a UK-CCC facility which is supported by HEFCE and PPARC.

REFERENCES

- Adelberger K. L., Shapley A. E., Steidel C. C., Pettini M., Erb D. K., Reddy N. A., 2005, *ApJ*, 629, 636
- Adelberger K. L., Steidel C. C., Shapley A. E., Pettini M., 2003, *ApJ*, 584, 45
- Aguirre A., Hernquist L., Katz N., Gardner J., Weinberg D., 2001, *ApJ*, 556, L11
- Aguirre A., Hernquist L., Schaye J., Weinberg D. H., Katz N., Gardner J., 2001, *ApJ*, 560, 599
- Aguirre A., Schaye J., Kim T.-S., Theuns T., Rauch M., Sargent W. L. W., 2004, *ApJ*, 602, 38
- Astier P., et al., 2006, *A&A*, 447, 31
- Balbi A., Quercellini C., 2007, *MNRAS*, 382, 1623
- Barkana R., Loeb A., 1999, *ApJ*, 523, 54
- Bassett B. A., Corasaniti P. S., Kunz M., 2004, *ApJ*, 617, L1
- Bechtold J., Crotts A. P. S., Duncan R. C., Fang Y., 1994, *ApJ*, 437, L83
- Becker G. D., Sargent W. L. W., Rauch M., 2004, *ApJ*, 613, 61
- Bergeron J., Boissé P., 1991, *A&A*, 243, 344
- Bertone S., Stoehr F., White S. D. M., 2005, *MNRAS*, 359, 1201
- Bi H., Davidsen A. F., 1997, *ApJ*, 479, 523
- Bolton J. S., Haehnelt M. G., Viel M., Springel V., 2005, *MNRAS*, 357, 1178
- Bouché N., Murphy M. T., Péroux C., 2004, *MNRAS*, 354, L25

- Bouché N., Murphy M. T., Péroux C., Csabai I., Wild V., 2006, MNRAS, 371, 495
- Bouchy F., Pepe F., Queloz D., 2001, A&A, 374, 733
- Caldwell R. R., 2002, Phys. Lett. B, 545, 23
- Caldwell R. R., Dave R., Steinhardt P. J., 1998, Phys. Rev. Lett., 80, 1582
- Caldwell R. R., Kamionkowski M., 2004, J. Cosmol. Astro-Part. Phys., 9, 9
- Capozziello S., Cardone V. F., Carloni S., DeMartino S., Falanga M., Troisi A., Bruni M., 2005, J. Cosmol. Astro-Part. Phys., 4, 5
- Carroll S. M., de Felice A., Duvvuri V., Easson D. A., Trodden M., Turner M. S., 2005, Phys. Rev. D, 71, 063513
- Carroll S. M., Press W. H., Turner E. L., 1992, ARA&A, 30, 499
- Cen R., Miralda-Escudé J., Ostriker J. P., Rauch M., 1994, ApJ, 437, L9
- Cen R., Nagamine K., Ostriker J. P., 2005, ApJ, 635, 86
- Cen R., Simcoe R. A., 1997, ApJ, 483, 8
- Charlton J. C., Anninos P., Zhang Y., Norman M. L., 1997, ApJ, 485, 26
- Charlton J. C., Churchill C. W., Linder S. M., 1995, ApJ, 452, L81
- Churchill C. W., Kacprzak G. G., Steidel C. C., Evans J. L., 2007, ApJ, 661, 714
- Churchill C. W., Mellon R. R., Charlton J. C., Jannuzi B. T., Kirhakos S., Steidel C. C., Schneider D. P., 2000, ApJ, 543, 577
- Churchill C. W., Steidel C. C., Vogt S. S., 1996, ApJ, 471, 164
- Cole S., et al., 2005, MNRAS, 362, 505
- Coppolani F., et al., 2006, MNRAS, 370, 1804
- Corasaniti P.-S., Huterer D., Melchiorri A., 2007, Phys. Rev. D, 75, 062001
- Cowie L. L., Songaila A., Kim T.-S., Hu E. M., 1995, AJ, 109, 1522
- Cristiani S., D'Odorico S., Fontana A., Giallongo E., Savaglio S., 1995, MNRAS, 273, 1016
- Curran S. J., Murphy M. T., Pihlström Y. M., Webb J. K., Bolatto A. D., Bower G. C., 2004, MNRAS, 352, 563
- Davis M. M., May L. S., 1978, ApJ, 219, 1
- Davis T. M., et al., 2007, ApJ, 666, 716
- Deffayet C., Dvali G., Gabadadze G., 2002, Phys. Rev. D, 65, 044023
- Desjacques V., Nusser A., Haehnelt M. G., Stoehr F., 2004, MNRAS, 350, 879
- Dinshaw N., Impey C. D., Foltz C. B., Weymann R. J., Chaffee F. H., 1994, ApJ, 437, L87
- D'Odorico V., Cristiani S., D'Odorico S., Fontana A., Giallongo E., Shaver P., 1998, A&A, 339, 678
- Dvali G., Gabadadze G., Porrati M., 2000, Phys. Lett. B, 485, 208
- Ebert R., Trümper M., 1975, in Bergmann P. G., Fenyves E. J., Motz L., eds, Seventh Texas Symposium on Relativistic Astrophysics Vol. 262 of Ann. N.Y. Acad. Sci. New York Academy of Sciences, New York, p. 470
- Eisenstein D. J., et al., 2005, ApJ, 633, 560
- Ellison S. L., Songaila A., Schaye J., Pettini M., 2000, AJ, 120, 1175
- Fan X., et al., 2006, AJ, 131, 1203
- Fernández-Soto A., Lanzetta K. M., Barcons X., Carswell R. F., Webb J. K., Yahil A., 1996, ApJ, 460, L85
- Francis P. J., 1993, ApJ, 407, 519
- Freedman W. L., et al., 2001, ApJ, 553, 47
- Freese K., Lewis M., 2002, Phys. Lett. B, 540, 1
- Gilmozzi R., Spyromilio J., 2007, The Messenger, 127, 11
- Gnedin N. Y., 1998, MNRAS, 294, 407
- Gnedin N. Y., Ostriker J. P., 1997, ApJ, 486, 581
- Gudmundsson E. H., Björnsson G., 2002, ApJ, 565, 1
- Haardt F., Madau P., 1996, ApJ, 461, 20
- Hambly N. C., et al., 2001, MNRAS, 326, 1279
- Heavens A., 2003, MNRAS, 343, 1327
- Hernquist L., Katz N., Weinberg D. H., Miralda-Escudé J., 1996, ApJ, 457, L51
- Hu E. M., Kim T.-S., Cowie L. L., Songaila A., Rauch M., 1995, AJ, 110, 1526
- Hubble E., 1929, Proc. Nat. Acad. Sci., 15, 168
- Hui L., Gnedin N. Y., 1997, MNRAS, 292, 27
- Hui L., Gnedin N. Y., Zhang Y., 1997, ApJ, 486, 599
- Jain B., Taylor A., 2003, Phys. Rev. Lett., 91, 141302
- John M. V., 2005, ApJ, 630, 667
- Kamenshchik A., Moschella U., Pasquier V., 2001, Phys. Lett. B, 511, 265
- Katz N., Weinberg D. H., Hernquist L., 1996, ApJS, 105, 19
- Kim T.-S., Carswell R. F., Cristiani S., D'Odorico S., Giallongo E., 2002, MNRAS, 335, 555
- Kim T.-S., Cristiani S., D'Odorico S., 2001, A&A, 373, 757
- Kim T.-S., Hu E. M., Cowie L. L., Songaila A., 1997, AJ, 114, 1
- Kirkman D., Tytler D., 1997, ApJ, 484, 672
- Lake K., 1981, ApJ, 247, 17
- Lake K., 1982, ApL, 23, 23
- Lake K., 2007, Phys. Rev. D, 76, 063508
- Linder E. V., 2003, Phys. Rev. Lett., 90, 091301
- Linder E. V., 2005, Phys. Rev. D, 72, 043529
- Linder E. V., Jenkins A., 2003, MNRAS, 346, 573
- Liske J., Webb J. K., Williger G. M., Fernández-Soto A., Carswell R. F., 2000, MNRAS, 311, 657
- Loeb A., 1998, ApJ, 499, L111
- Lovis C., et al., 2005, A&A, 437, 1121
- Lu L., Sargent W. L. W., Womble D. S., Takada-Hidai M., 1996, ApJ, 472, 509
- McVittie G. C., 1962, ApJ, 136, 334
- McVittie G. C., 1965, General Relativity and Cosmology. University of Illinois Press, Urbana
- Madau P., Ferrara A., Rees M. J., 2001, ApJ, 555, 92
- Maor I., Brustein R., Steinhardt P. J., 2001, Phys. Rev. Lett., 86, 6
- Miralda-Escudé J., Cen R., Ostriker J. P., Rauch M., 1996, ApJ, 471, 582
- Murphy M. T., et al., 2007, MNRAS, 380, 839
- Nakamura T., Chiba T., 1999, MNRAS, 306, 696
- Oke J. B., 1974, ApJS, 27, 21
- Overduin J. M., Cooperstock F. I., 1998, Phys. Rev. D, 58, 043506
- Partovi M. H., Mashhoon B., 1984, ApJ, 276, 4
- Pasquini L., et al., 2005, ESO publication OWL-CSR-ESO-00000-0160
- Pasquini L., et al., 2006, in Herbst T., ed., Instrumentation for Extremely Large Telescopes Vol. 0106 of MPIA Special Publication. MPIA, Heidelberg, p. 127
- Peacock J. A., 1999, Cosmological Physics. Cambridge University Press, Cambridge
- Peacock J. A., et al., 2001, Nat, 410, 169
- Peebles P. J. E., 1993, Principles of Physical Cosmology. Princeton University Press, Princeton, New Jersey
- Perlmutter S., et al., 1999, ApJ, 517, 565
- Pettini M., Shapley A. E., Steidel C. C., Cuby J.-G., Dickinson M., Moorwood A. F. M., Adelberger K. L., Giallisco M., 2001, ApJ, 554, 981
- Phillipps S., 1982, ApL, 22, 123
- Pieri M. M., Haehnelt M. G., 2004, MNRAS, 347, 985
- Primack J. R., 2005, NewAR, 49, 25

- Prochaska J. X., Chen H.-W., Wolfe A. M., Dessauges-Zavadsky M., Bloom J. S., 2008, *ApJ*, 672, 59
- Ratra B., Peebles P. J. E., 1988, *Phys. Rev. D*, 37, 3406
- Rauch M., 1998, *ARA&A*, 36, 267
- Rauch M., Becker G. D., Viel M., Sargent W. L. W., Smette A., Simcoe R. A., Barlow T. A., Haehnelt M. G., 2005, *ApJ*, 632, 58
- Rauch M., Haehnelt M. G., 1995, *MNRAS*, 275, L76
- Rauch M., Sargent W. L. W., Barlow T. A., Carswell R. F., 2001, *ApJ*, 562, 76
- Richards G. T., et al., 2002, *AJ*, 123, 2945
- Riess A. G., et al., 1998, *AJ*, 116, 1009
- Riess A. G., et al., 2004, *ApJ*, 607, 665
- Rollinde E., Petitjean P., Pichon C., Colombi S., Aracil B., D’Odorico V., Haehnelt M. G., 2003, *MNRAS*, 341, 1279
- Rüdiger R., 1980, *ApJ*, 240, 384
- Rüdiger R., 1982, *ApJ*, 260, 33
- Sandage A., 1962, *ApJ*, 136, 319
- Saslaw W. C., 1989, *ApJ*, 341, 588
- Schaye J., 2001, *ApJ*, 559, 507
- Schaye J., Aguirre A., Kim T.-S., Theuns T., Rauch M., Sargent W. L. W., 2003, *ApJ*, 596, 768
- Schaye J., Rauch M., Sargent W. L. W., Kim T.-S., 2000, *ApJ*, 541, L1
- Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
- Schneider D. P., et al., 2007, *AJ*, 134, 102
- Seljak U., Slosar A., McDonald P., 2006, *J. Cosmol. Astro-Part. Phys.*, 10, 14
- Seljak U., Zaldarriaga M., 1996, *ApJ*, 469, 437
- Seo H.-J., Eisenstein D. J., 2003, *ApJ*, 598, 720
- Seto N., Cooray A., 2006, *Phys. Rev. D*, 73, 023005
- Shapiro C., Turner M. S., 2006, *ApJ*, 649, 563
- Simcoe R. A., Sargent W. L. W., Rauch M., 2004, *ApJ*, 606, 92
- Simcoe R. A., Sargent W. L. W., Rauch M., Becker G., 2006, *ApJ*, 637, 648
- Smette A., Robertson J. G., Shaver P. A., Reimers D., Wisotzki L., Koehler T., 1995, *A&AS*, 113, 199
- Songaila A., 2006, *AJ*, 131, 24
- Spergel D. N., et al., 2003, *ApJS*, 148, 175
- Spergel D. N., et al., 2007, *ApJS*, 170, 377
- Springel V., 2005, *MNRAS*, 364, 1105
- Steidel C. C., 1995, in Meylan G., ed., *QSO Absorption Lines* Springer, Berlin, p. 139
- Steidel C. C., Dickinson M., Persson S. E., 1994, *ApJ*, 437, L75
- Szydlowski M., Kurek A., Krawiec A., 2006, *Phys. Lett. B*, 642, 171
- Teuber J., 1986, *ApL*, 25, 139
- Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, *MNRAS*, 301, 478
- Theuns T., Mo H. J., Schaye J., 2001, *MNRAS*, 321, 450
- Theuns T., Schaye J., Haehnelt M. G., 2000, *MNRAS*, 315, 600
- Theuns T., Viel M., Kay S., Schaye J., Carswell R. F., Tzanavaris P., 2002, *ApJ*, 578, L5
- Turner M. S., White M., 1997, *Phys. Rev. D*, 56, 4439
- Véron-Cetty M.-P., Véron P., 2006, *A&A*, 455, 773
- Viel M., Haehnelt M. G., Lewis A., 2006, *MNRAS*, 370, L51
- Viel M., Haehnelt M. G., Springel V., 2004, *MNRAS*, 354, 684
- Viel M., Matarrese S., Mo H. J., Theuns T., Haehnelt M. G., 2002, *MNRAS*, 336, 685
- Wang Y., 2006, *ApJ*, 647, 1
- Wang Y., Freese K., Gondolo P., Lewis M., 2003, *ApJ*, 594, 25
- Wang Y., Tegmark M., 2005, *Phys. Rev. D*, 71, 103513
- Weinberg S., 1972, *Gravitation and Cosmolgy*. John Wiley & Sons, New York
- White S. D. M., 2007, *Rep. Prog. Phys.*, 70, 883
- Wiklind T., Combes F., 1999, in Carilli C. L., Radford S. J. E., Menten K. M., Langston G. I., eds, *Highly Redshifted Radio Lines Vol. 156 of ASP Conf. Ser. Astronomical Society of the Pacific*, San Francisco, p. 202
- Zhang H., Zhong W., Zhu Z.-H., He S., 2007, *Phys. Rev. D*, 76, 123508
- Zhang Y., Anninos P., Norman M. L., 1995, *ApJ*, 453, L57
- Zhang Y., Meiksin A., Anninos P., Norman M. L., 1998, *ApJ*, 495, 63
- Zhu Z.-H., Fujimoto M.-K., 2004, *ApJ*, 602, 12
- Zibetti S., Ménard B., Nestor D., Turnshek D., 2005, *ApJ*, 631, L105

APPENDIX A: REDSHIFT DRIFT INCLUDING PECULIAR MOTION

Equation (4) for the redshift drift is only correct if both the observer and the source are comoving. In reality, both will have peculiar velocities and accelerations due to local gravitational potential wells. The purpose of this Appendix is to provide an exact calculation of the effects of peculiar motion on the observed redshift drift.

Our own peculiar motion can be broken down into several components, such as the Earth’s motion with respect to the solar system’s barycentre, the solar system’s rotation within the Galaxy, and the Galaxy’s movement with respect to the Hubble flow. All of these components are or will soon be known with high precision so that the observations can be corrected accordingly (Pasquini et al. 2006). Hence we will not consider the observer’s peculiar motion any further.

For a given object at large distance it is in general not possible to determine its peculiar motion or to correct its observed redshift for it. However, the peculiar motions of widely separated objects are expected to be uncorrelated. Hence, when using several such objects to measure \dot{z} , their peculiar motions are only expected to introduce a random noise component, but will not cause any systematic bias. To assess the relevance of peculiar motions for any given implementation of a \dot{z} experiment one must compare the noise induced by the peculiar motions of the experiment’s targets with other sources of error (ideally only photon noise).

We now derive the expression for the observed redshift drift in the presence of peculiar velocity and acceleration. Consider a photon emitted at time t_{em} by a source moving with velocity v with respect to its local Hubble flow (object A in Fig. A1), and received by a comoving observer at time t_{obs} (object B in Fig. A1). Clearly, the observed redshift of this photon travelling from A to B is given by

$$1 + z_{\text{obs}}(t_{\text{obs}}, t_{\text{em}}) = [1 + z_{\text{D}}(t_{\text{em}})] [1 + z(t_{\text{obs}}, t_{\text{em}})]. \quad (\text{A1})$$

Here, z is the cosmological redshift as before, and z_{D} is the (special relativistic) Doppler shift of the emitter:

$$1 + z_{\text{D}}(t_{\text{em}}) = \frac{1 + \cos[\theta(t_{\text{em}})] \frac{v(t_{\text{em}})}{c}}{\sqrt{1 - [\frac{v(t_{\text{em}})}{c}]^2}} = \gamma (1 + \cos \theta \frac{v}{c}), \quad (\text{A2})$$

where θ is the angle between the emitter’s peculiar velocity and the outward line of sight (LOS) from the observer (see Fig. A1). The redshift drift is observed as the redshift difference between the




Figure A1. Cartoon illustrating various photons, their emitters and observers, as well as their times of emission and observation (see text for details). We also show various angles that define the peculiar motion of object A: the angle θ in panel (a) is the angle between the peculiar velocity \mathbf{v} and the LOS, while η is the angle between \mathbf{v} and the peculiar acceleration projected into the plane of the drawing, \mathbf{a}_p . The acceleration \mathbf{a} itself makes an angle ζ with this plane (not shown). In panel (b) \mathbf{v} has acquired a component perpendicular to the plane of the drawing (making an angle ξ), and so we only show its projected component, \mathbf{v}_p , as well as its angle θ_p with the LOS. Compared to panel (a) \mathbf{v} has shifted by an angle ψ in the plane of the drawing while the direction of the LOS has changed by the angle ϕ .

photons labelled as 1 and 2 in Fig. A1, and has two terms:

$$\frac{z_{\text{obs}}(t_{\text{obs}}^2, t_{\text{em}}^2) - z_{\text{obs}}(t_{\text{obs}}^1, t_{\text{em}}^1)}{t_{\text{obs}}^2 - t_{\text{obs}}^1} \approx \frac{dz_{\text{obs}}}{dt_{\text{obs}}} = \frac{dz_D}{dt_{\text{obs}}} (1+z) + (1+z_D) \frac{dz}{dt_{\text{obs}}}. \quad (\text{A3})$$

We can therefore identify the following, separable effects of peculiar motions on the process of putting data points on Fig. 2: (i) The Doppler shift in equation (A1) introduces an uncertainty on where to place a given object along the x -axis of Fig. 2. Uncertainties in the placement along the y -axis are caused by (ii) peculiar acceleration, (iii) the Doppler shift in the second term of equation (A3) and (iv) the change in the comoving coordinates of the source. This last point can in turn be separated into two distinct effects: the decrease of θ (due to parallax) and the change of the distance between the observer and the source. The latter implies that the dz/dt_{obs} appearing in equation (A3) is *not* the same as that of equation (4), which is only valid for a constant χ .

This is illustrated by photons 3 and 4 in Fig. A1. Photon 3 is emitted at t_{em}^1 by the *comoving* object C that is at the same location as the moving object A at time t_{em}^1 , while photon 4 is emitted by comoving object D just as A is passing by at time t_{em}^2 . Photons 3 and 4 are received by B at the same times as photons 1 and 2, respectively (i.e. t_{obs}^1 and t_{obs}^2). Clearly, the redshift difference between these two photons is the cosmological shift we are seeking:

$$dz = z_{D \rightarrow B}(t_{\text{obs}}^2, t_{\text{em}}^2) - z_{C \rightarrow B}(t_{\text{obs}}^1, t_{\text{em}}^1) \quad (\text{A4})$$

To calculate dz we make use of photon 5 in Fig. A1. Like photon 3 this photon is sent from C to B, although at a different time. Since C and B are comoving objects, the redshift difference of photons 3 and 5 is given by equation (4). We have:

$$dz = z_{D \rightarrow B}(t_{\text{obs}}^2, t_{\text{em}}^2) - z_{C \rightarrow B}(t_{\text{obs}}^2, t_{\text{em}}^2 + \Delta t_{\text{em}}) + z_{C \rightarrow B}(t_{\text{obs}}^2, t_{\text{em}}^2 + \Delta t_{\text{em}}) - z_{C \rightarrow B}(t_{\text{obs}}^1, t_{\text{em}}^1)$$

$$\begin{aligned} &\approx -\frac{\partial z}{\partial t_{\text{em}}} \Delta t_{\text{em}} + \frac{dz_{\chi}}{dt_{\text{obs}}} dt_{\text{obs}} \\ &= (1+z) H(t_{\text{em}}) \Delta t_{\text{em}} \\ &\quad + [(1+z) H(t_{\text{obs}}) - H(t_{\text{em}})] dt_{\text{obs}} \\ &= (1+z) H(t_{\text{obs}}) dt_{\text{obs}} \\ &\quad - (1+z) H(t_{\text{em}}) (dt_C - \Delta t_{\text{em}}) \\ &= (1+z) H(t_{\text{obs}}) (t_{\text{obs}}^2 - t_{\text{obs}}^1) \\ &\quad - (1+z) H(t_{\text{em}}) (t_{\text{em}}^2 - t_{\text{em}}^1), \end{aligned} \quad (\text{A5})$$

where Δt_{em} is the time required by photon 4 to cover the comoving radial coordinate difference, $\Delta\chi$, between C and D as seen from B. dt_C is the time difference between the emission of photons 3 and 5 by C, so that $dt_{\text{obs}} = (1+z) dt_C$. The last equality of equation (A5) holds by construction. Using

$$c\Delta t_{\text{em}} = a(t_{\text{em}})\Delta\chi = \cos\theta v (t_{\text{em}}^2 - t_{\text{em}}^1) \quad (\text{A6})$$

one obtains

$$\frac{t_{\text{em}}^2 - t_{\text{em}}^1}{t_{\text{obs}}^2 - t_{\text{obs}}^1} = [(1+z) (1 + \cos\theta \frac{v}{c})]^{-1} \quad (\text{A7})$$

and hence

$$\begin{aligned} \frac{dz}{dt_{\text{obs}}} &= (1+z) H(t_{\text{obs}}) - (1 + \cos\theta \frac{v}{c})^{-1} H(t_{\text{em}}) \\ &= (1+z) H(t_{\text{obs}}) - \frac{\gamma}{1+z_D} H(t_{\text{em}}). \end{aligned} \quad (\text{A8})$$

Turning now to the first term of equation (A3) the change in the Doppler shift is given by:

$$dz_D = \frac{dz_D}{dt_{\text{em}}} (t_{\text{em}}^2 - t_{\text{em}}^1) \quad (\text{A9})$$

so that

$$\frac{dz_D}{dt_{\text{obs}}} = \frac{dz_D}{dt_{\text{em}}} [(1+z) (1 + \cos\theta \frac{v}{c})]^{-1}, \quad (\text{A10})$$

where dz_D/dt_{em} is the rate of change of the Doppler shift as measured by a local comoving observer along the LOS to B. Using a prime to denote differentiation with respect to t_{em} , we have

$$z'_D = \frac{v'}{c} (1+z_D) \left(\frac{\cos\theta}{1 + \cos\theta \frac{v}{c}} + \gamma^2 \frac{v}{c} \right) - \gamma \sin\theta \frac{v}{c} \theta'. \quad (\text{A11})$$

In general, the peculiar acceleration, \mathbf{a} , can have any orientation with respect to \mathbf{v} . We denote the direction of \mathbf{a} by the angles ζ and η . The former is the angle between \mathbf{a} and the plane spanned by \mathbf{v} and the observer (i.e. the plane of Fig. A1). The latter is the angle between \mathbf{a}_p , the projection of \mathbf{a} into this plane, and \mathbf{v} [see Fig. A1(a)]. Hence we have

$$v' = \cos\eta a_p = \cos\eta \cos\zeta a. \quad (\text{A12})$$

The angle θ changes with time because both the viewing angle and the direction of \mathbf{v} change. In particular, \mathbf{v} may acquire a component that is perpendicular to the plane of Fig. A1, so that it makes some angle ξ with this plane. Writing θ_p for the angle between the LOS and the projected velocity \mathbf{v}_p (see Fig. A1), we have

$$\cos\theta(t_{\text{em}}) = \cos\theta_p(t_{\text{em}}) \cos\xi(t_{\text{em}}) \quad (\text{A13})$$

and hence

$$\sin\theta \theta' = \cos\xi \sin\theta_p \theta'_p + \cos\theta_p \sin\xi \xi'. \quad (\text{A14})$$

However, since this equation should be evaluated at t_{em}^1 and since $\xi(t_{\text{em}}^1) = 0$ and $\theta(t_{\text{em}}^1) = \theta_p(t_{\text{em}}^1)$ by construction, we can simply ignore the perpendicular component and write $\theta'(t_{\text{em}}^1) = \theta'_p(t_{\text{em}}^1)$

[unless $\theta(t_{\text{em}}^1) = 0$, in which case the second term of equation (A11) is zero anyway].

The change in θ_p is given by the angles ϕ and ψ between the dashed and solid lines in Fig. A1(b). Clearly, ϕ is the parallax angle for a baseline of $\sin \theta v (t_{\text{em}}^2 - t_{\text{em}}^1)$, i.e. the transverse distance between C and D, while ψ denotes the change of v 's direction in the plane of the drawing. Hence we obtain

$$\begin{aligned} \theta' &= \frac{\psi - \phi}{t_{\text{em}}^2 - t_{\text{em}}^1} = \sin \eta \cos \zeta \frac{a}{v} - \frac{\sin \theta v}{d_{\text{px}}(t_{\text{em}}, \chi)} \\ &= \sin \eta \cos \zeta \frac{a}{v} - \frac{1+z}{d_{\text{px}}(t_{\text{obs}}, \chi)} \sin \theta v. \end{aligned} \quad (\text{A15})$$

d_{px} is the parallax distance between two objects separated by a co-moving coordinate distance χ (e.g. Peebles 1993; Peacock 1999):

$$d_{\text{px}}(t, \chi) = a(t) \frac{\Sigma(\chi)}{\Xi(\chi)}, \quad (\text{A16})$$

where we have used the notation

$$\Sigma(\chi) = \begin{cases} \sin \chi & k = +1 \\ \chi & k = 0 \\ \sinh \chi & k = -1 \end{cases} \quad (\text{A17})$$

and its cosine equivalent $\Xi(\chi) = \sqrt{1 - k\Sigma^2(\chi)}$.

Collating all of the above we finally find for the observed redshift drift:

$$\begin{aligned} \frac{dz_{\text{obs}}}{dt_{\text{obs}}} &= \gamma \cos \zeta \cos \eta \frac{a}{c} \left(\frac{\cos \theta}{1 + \cos \theta \frac{v}{c}} + \gamma^2 \frac{v}{c} \right) \\ &\quad - \gamma \cos \zeta \sin \eta \frac{a}{c} \frac{\sin \theta}{1 + \cos \theta \frac{v}{c}} \\ &\quad + \gamma \frac{1+z}{d_{\text{px}}} \frac{\sin^2 \theta \frac{v^2}{c}}{1 + \cos \theta \frac{v}{c}} \\ &\quad + (1 + z_{\text{obs}}) H(t_{\text{obs}}) - \gamma H(t_{\text{em}}) \\ &= (1 + z_{\text{obs}}) H_0 \\ &\quad - \gamma \left\{ H(z) - \cos \zeta \frac{a}{c} \left[\frac{\cos(\theta + \eta)}{1 + \cos \theta \frac{v}{c}} + \cos \eta \gamma^2 \frac{v}{c} \right] \right\} \\ &\quad + \frac{1 + z_{\text{obs}}}{d_{\text{px}}(z)} \frac{v_t^2}{c} \\ &= (1 + z_{\text{obs}}) H_0 - \gamma \left[H(z) - \frac{a_r/c}{1 + \cos \theta \frac{v}{c}} - \gamma^2 \frac{vv'}{c^2} \right] \\ &\quad + \frac{1 + z_{\text{obs}}}{d_{\text{px}}(z)} \frac{v_t^2}{c} \end{aligned} \quad (\text{A18})$$

where we have used the notation

$$v_t = \frac{\sin \theta v}{1 + \cos \theta \frac{v}{c}} \quad (\text{A19})$$

for the emitter's apparent transverse velocity and

$$a_r = \cos \zeta \cos(\theta + \eta) a \quad (\text{A20})$$

for the magnitude of a 's component along the LOS. Comparing with equation (4) we can see that the main difference is the modification of $H(z)$ with a term proportional to the change in the radial velocity. The second order term stems from taking the derivative of γ in equation (A11), while the last term is due to the parallax effect. Clearly, equation (A18) reduces to equation (4) for $v = a = 0$. Note also that ζ is entirely degenerate with a and so we will set $\zeta = 0$ from now on.

The formulae for many special cases such as $a = 0$ (constant motion), $\theta = 0$ (radial motion) and $\theta = \pi/2$ (transverse

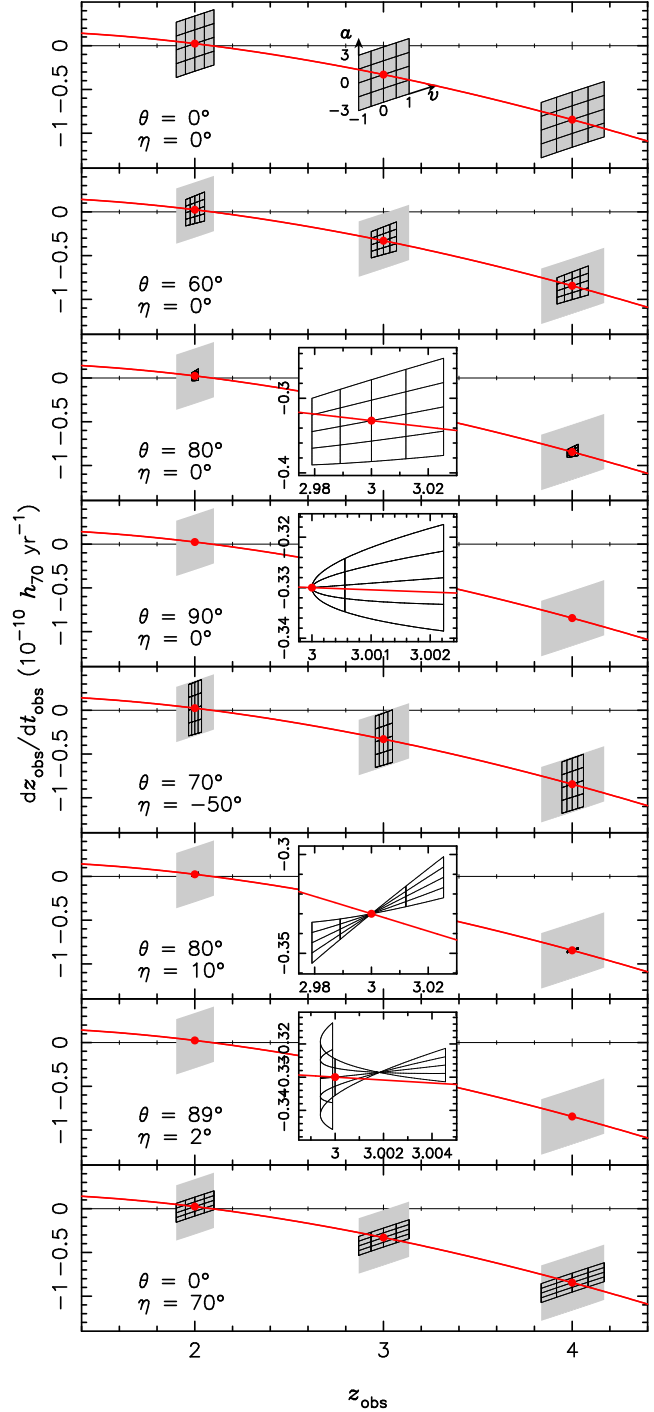


Figure A2. Top panel: the solid line shows the cosmological redshift drift, $\dot{z}(z)$, expected for comoving sources (cf. Fig. 2). The three grids show how an object at $z = 2, 3$ and 4 , respectively, moves in the observed $z_{\text{obs}} - \dot{z}_{\text{obs}}$ plane as a function of v and a , assuming $\theta = \eta = \zeta = 0$. For each grid, the vertical lines are lines of constant v , while the tilted horizontal lines are lines of constant a . The values of v and a are indicated along the middle grid in units of 10^4 km s^{-1} and $10^{-8} \text{ cm s}^{-2}$. The other panels show the same for different values of θ and η as indicated (and $\zeta = 0$). For comparison, the grey shaded areas in each panel are the same as in the top panel. The insets show a blow-up of the region around the $z = 3$ point.

motion) are immediately apparent from equation (A18). Another special case is that of a freely moving source. In the absence of local gravitational potential wells we have $v \propto a^{-1}$ (and $\zeta = \eta = 0$) so that $a = v' = -vH(t_{\text{em}})$. Neglecting terms of order $(\frac{v}{c})^4$ and higher, we obtain

$$\frac{\dot{z}_{\text{obs}}}{1+z_{\text{obs}}} \approx H_0 - \frac{H(z)}{1+z} \left(1 + \frac{v_t^2}{c^2}\right) + \frac{1}{d_{\text{px}}(z)} \frac{v_t^2}{c}, \quad (\text{A21})$$

which shows that in this case purely radial peculiar motion preserves the form of equation (4). Setting $\Omega_\Lambda = 0$, we recover the formula of Teuber (1986):

$$\frac{\dot{z}_{\text{obs}}}{1+z_{\text{obs}}} \approx H_0 \left[1 - \sqrt{1 + \Omega_M z} \left(1 + \frac{v_t^2}{c^2}\right)\right] + \frac{1}{d_{\text{px}}(z)} \frac{v_t^2}{c}. \quad (\text{A22})$$

Given an object's cosmological redshift z , and the parameters of its peculiar motion (v , θ) and (a , η) we can now predict its observed position in a \dot{z}_{obs} vs. z_{obs} plot, i.e. the observed equivalent of Fig. 2. This is shown in Fig. A2 for three objects at $z = 2, 3$ and 4 , respectively. The solid line in each panel shows the cosmological $\dot{z}(z)$ that one expects in the absence of peculiar motions. In each panel we fix θ and η at some value and use small grids to show how the three objects are perturbed away from their 'correct' positions (marked as dots) as a function of v and a . The grids cover the range -1 to $+1 \times 10^4 \text{ km s}^{-1}$ in v and -3 to $+3 \times 10^{-8} \text{ cm s}^{-2}$ in a . To relate these scales to one another we point out that a constant acceleration of $1.5 \times 10^{-8} \text{ cm s}^{-2}$ would result in a peculiar velocity of 10^4 km s^{-1} by $z = 3$.

Obviously, the largest possible effects are obtained when $\theta = \eta = 0$ (top panel), and we reproduce this case as grey shaded areas in all other panels for comparison. The extent of the grids in the v and a directions can be scaled independently and almost arbitrarily by appropriately adjusting θ and η . However, when $\theta + \eta \approx \pi/2$ purely relativistic and/or second order effects take over. These are a factor of ~ 10 – 100 smaller in magnitude and can significantly distort the grids. Note that since the grids are generally not symmetric, it is in principle possible (contrary to our initial expectation) that peculiar motions could bias a $\dot{z}(z)$ measurement, even when using a large sample of objects whose peculiar motions are uncorrelated.

Given typical values for the peculiar velocities and accelerations of a particular set of objects (e.g. Figs. 3 and 4), we can essentially read off Fig. A2 the scale of the noise on a $\dot{z}(z)$ measurement introduced by the peculiar motions. This must be compared to the photon noise on $\dot{z}(z)$ for an *individual* object. In the case of the Ly α forest, as discussed in this paper, the photon noise error on \dot{z} for an individual absorption line is about four orders of magnitude larger than the error induced by peculiar motions of the absorbing gas. We conclude that peculiar motions will only be relevant for a \dot{z} experiment if it is based on very precise measurements on a small number of objects with large peculiar velocities and accelerations.

APPENDIX B: SUMMARY OF SELECTED ALTERNATIVE COSMOLOGICAL MODELS

There exist a number of cosmological models that may offer an alternative to the standard Λ CDM scenario. The purpose of this Appendix is to provide a summary of the predictions for the redshift drift of an incomplete subset of these alternative models. All of them were introduced specifically to explain the observed acceleration of the universal expansion, and they can be separated into two distinct classes: (i) models that modify the stress-energy tensor of the Universe to include a new dark energy component, but

leave the general relativistic field equation intact and (ii) models that modify the theory of gravity.

Recall that the validity of the Robertson-Walker metric was the only assumption we needed to derive equation (4) for the redshift drift (cf. Section 2). None of the models considered in this Appendix violate this assumption, and so all we need to do in order to determine a given model's prediction for \dot{z} is to specify the model's Friedman equation (see also Szydlowski, Kurek & Krawiec 2006 for an overview).

B1 Alternative dark energy models

The general relativistic Friedman equation as given in equation (5) assumes a specific form for each mass-energy component's equation of state, namely $p_i = w_i c^2 \rho_i$, where p_i and ρ_i are the component's pressure and density, respectively. Without this assumption the Friedman equation reads:

$$\left[\frac{H(z)}{H_0}\right]^2 = \sum_i \Omega_i \frac{\rho_i(z)}{\rho_{i0}} + \Omega_k (1+z)^2, \quad (\text{B1})$$

where ρ_{i0} is the present-day value of ρ_i . For an arbitrary equation of state, $p_i(\rho_i)$, the evolution of ρ_i is obtained by solving the general relativistic energy conservation equation:

$$\frac{d\rho_i}{da} = -\frac{3}{a} \left[\rho_i + \frac{p_i(\rho_i)}{c^2} \right]. \quad (\text{B2})$$

B1.1 Quintessence

Here, we take the term quintessence to simply mean a smooth, time varying dark energy component with equation of state $p_X = w_X c^2 \rho_X$, where w_X may also vary with time. In principle, w_X and its evolution should be given by the scalar field equation of the underlying physical theory of this component. In practice, however, it is common to use a simple parameterisation of w_X in order to be less model dependent. Here we will use the expression of Linder (2003):

$$w_X(z) = w_0 + w_a \left(1 - \frac{a}{a_0}\right) = w_0 + w_a \frac{z}{1+z}. \quad (\text{B3})$$

In this case the Friedman equation is given by

$$\left[\frac{H(z)}{H_0}\right]^2 = \Omega_X (1+z)^{3(1+w_0+w_a)} e^{-3w_a z/(1+z)} + \Omega_M (1+z)^3 + \Omega_k (1+z)^2, \quad (\text{B4})$$

where we have included a matter component with $w_M = 0$. Note that for $(w_0, w_a) = (-1, 0)$ the above reduces to the standard Λ CDM case.

B1.2 Generalised Chaplygin gas

In the Λ CDM and quintessence models it is the different evolution of the densities of the two mass-energy components that causes the transition from a decelerating (matter-dominated) to an accelerating (Λ /quintessence-dominated) Universe. In contrast, the generalised Chaplygin gas (GCG) model seeks to explain this transition with only a *single* background fluid. This leads to a rather exotic equation of state for GCG:

$$p_c = -A c^2 \rho_c^{-\alpha}, \quad (\text{B5})$$

where A and α are constants with $A \geq 0$ and $\alpha \geq -1$. Clearly, for $\alpha = -1$ GCG is identical to Λ and therefore the model is identical to the de Sitter Universe (containing Λ but no matter), while $\alpha = 1$ corresponds to the standard Chaplygin gas model (Kamenshchik et al. 2001). The above equation of state results in the following Friedman equation:

$$\left[\frac{H(z)}{H_0} \right]^2 = \Omega_c [B + (1 - B)(1 + z)^{3(1+\alpha)}]^{1/(1+\alpha)} + \Omega_k (1 + z)^2, \quad (\text{B6})$$

where $B = A/\rho_{c0}^{1+\alpha}$. Note that we are neglecting the baryons which had been subsumed in Ω_M previously. In this case GCG is the only mass-energy component and we have $\Omega_c = 1 - \Omega_k$. Inserting the solution for ρ_c into the equation of state we obtain:

$$p_c = \frac{-Bc^2}{B + (1 - B)(1 + z)^{3(1+\alpha)}} \rho_c \equiv -w_c(z) c^2 \rho_c. \quad (\text{B7})$$

From the above equations we can see clearly that for $0 < B < 1$ the Hubble expansion and the GCG equation of state smoothly change from being matter-like at early times to being Λ -like at late times (possibly in the future), where the parameters B and α jointly control when this transition occurs and how long it takes. For the extreme values of $B = w_c(0) = 0$ and 1 GCG behaves like pure matter and Λ , respectively, regardless of the value of α . Note that for $\alpha = 0$ the GCG model reduces to the standard Λ CDM case.

B2 Alternative gravity models

B2.1 Cardassian models

In Cardassian⁴ models the Universe is assumed to be flat and to contain only matter but no dark energy. The acceleration of the expansion is instead achieved by altering the Friedman equation. In the modified polytropic Cardassian models the Friedman equation reads (Wang et al. 2003):

$$\left[\frac{H(z)}{H_0} \right]^2 = \Omega_M (1 + z)^3 [1 + (\Omega_M^{-q} - 1)(1 + z)^{3q(n-1)}]^{1/q}, \quad (\text{B8})$$

where the parameters $q > 0$ and $n < 2/3$ are constant. The case $q = 1$ corresponds to the original Cardassian model (Freese & Lewis 2002) which is, however, as far as the Hubble expansion is concerned, exactly equivalent to a flat quintessence model with $w_X = n - 1$. Note that in Cardassian models spatial flatness is *not* equivalent to $\Omega_M = 1$.

At early times the Cardassian Friedman equation is dominated by the first term inside the brackets, resulting in normal, matter-dominated expansion. However, at late times the non-standard term dominates and gives rise to the acceleration. Such modifications to the Friedman equation may be motivated by self-interacting dark matter or by considering our observable Universe as a 3-dimensional brane embedded in a higher dimensional space-time.

B2.2 DGP model

Another brane-world scenario is the Dvali, Gabadadze & Porrati (DGP; 2000) model. In this (4+1)-dimensional model gravity can spread (or ‘leak’) to the flat, infinite extra spatial dimension at

distances larger than some cross-over scale r_c . On scales smaller than r_c an observer on the brane thus measures normal (3+1)-dimensional gravity, while on larger scales gravity weakens, causing accelerated expansion beyond r_c . If r_c is chosen large enough then it will only enter the Hubble horizon at late times so that the early evolution of the Universe remains unaltered compared to standard, matter-dominated general relativity, while at late times the leaking of gravity to the extra dimension accelerates the expansion without requiring any form of dark energy.

Defining an r_c -induced density parameter,

$$\Omega_{r_c} = \frac{c^2}{4r_c^2 H_0^2}, \quad (\text{B9})$$

we can write the Friedman equation of the DGP model as (Deffayet et al. 2002):

$$\left[\frac{H(z)}{H_0} \right]^2 = \left[\sqrt{\Omega_{r_c}} + \sqrt{\Omega_{r_c} + \Omega_M (1 + z)^3} \right]^2 + \Omega_k (1 + z)^2. \quad (\text{B10})$$

Note that in this model the curvature parameter is given by:

$$\Omega_k = 1 - \left(\sqrt{\Omega_{r_c}} + \sqrt{\Omega_{r_c} + \Omega_M} \right)^2, \quad (\text{B11})$$

so that flatness implies

$$\Omega_{r_c} = \left(\frac{1 - \Omega_M}{2} \right)^2. \quad (\text{B12})$$

⁴ With dark energy being ‘bad’ (White 2007) and Cardassians being ‘ugly’, all we need now is a ‘good’ theory to complete the famous trio.